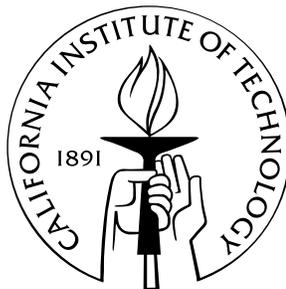


Neural Dynamics and Population Coding in the Insect Brain

Thesis by
OFER MAZOR

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2005
(Defended May 24, 2005)

© 2005

Ofer Mazor

All Rights Reserved

To my parents

Acknowledgements

I could not have completed the work presented in this thesis without the help of many people. I would like to thank my advisor, Gilles Laurent, for all his help and ideas, and for letting me explore the questions that interested me. I am grateful to the members of my thesis committee, Thanos Siapas, Pietro Perona, Michael Dickinson and Mark Konishi, who provided me with many helpful ideas and suggestions.

Significant portions of the work described in this thesis were the result of collaborations. I enjoyed working with Christophe Pouzat, Javier Perez-Orive, Stijn Cassenaer, Glenn Turner, and Rachel Wilson, without whom this work would not have been possible.

In addition, for my first three years at Caltech I was generously supported by a Department of Defense NDSEG Fellowship.



I am grateful to all the members of the Laurent Lab, who made the last six years of my life so enjoyable and intellectually stimulating. I would especially like to thank Christophe Pouzat, who is a great friend and was instrumental in my development as a scientist. He taught me to be skeptical, and he introduced me to cycling and jazz. I am indebted to Stijn Cassenaer for his friendship, generosity and wit, both in the lab and on the bike. The discussions I have had with Vivek Jayaraman over the past few years, about science and everything else, have been invaluable. I would also like to thank Roni Jortner, Dan Gold, Ben Rubin, Mark Stopfer, Sarah Farivar, Javier Perez-Orive, Maria Papadopoulou, Anusha Narayan, Cindy Chiu, Laurent Moreaux, Ania Mitros, Adam Hayes, and the rest of my first-year CNS class for contributing so much to my time at Caltech.



Most importantly, I must thank my family: Mommy, Abba, Yuval and Tali. Thank you for supporting me in everything I have done, even if you did tease me that I would never graduate.

Abstract

Sensory information is represented in the brain through the activity of populations of neurons. How this information is encoded and how it is processed and read out are crucial questions in neuroscience. The work presented here examines these issues using an insect brain model system. Specifically, this work addresses how odor information is represented across a population of neurons in this relatively simple nervous system. It asks how the dynamics of a population of neurons contribute to the encoding of information.

To address these questions, simultaneous multi-unit extracellular recordings were made in vivo in the locust brain. The first part of the dissertation describes several advances in spike-sorting methods that were necessary for analyzing such recordings. These advances include quantitative tests of sorting quality, and they allow for automated spike-sorting. Using these techniques, data sampled from tens of neurons over hours of recording can be analyzed with relative ease.

The remainder of the dissertation examines the encoding of olfactory information by a population of neurons called projection neurons (PNs), located in the first olfactory relay of the brain. Odor information is shown to be represented by a subpopulation of responsive PNs. The composition of this population changes

over time in an odor-specific manner, thus forming a distributed, dynamical representation. The statistics of this response and its dynamics are quantified.

Furthermore, the mechanism by which odor information is extracted from the PN population response is examined. A second set of recordings were made from Kenyon cells (KCs), which receive direct excitatory synaptic input from PNs. The dynamic response of the PN population appears to be decoded by KCs through a mechanism based on several underlying components, including oscillatory dynamics, feed-forward inhibition, and intrinsic properties of the KCs. This decoding process is shown to drastically change the odor representations, from dense to sparse.

Taken together, the results presented in this dissertation establish that the complex spatial and temporal dynamics of the PN population do encode odor information, and that this information is decoded by other neurons (KCs) in a very precise way, resulting in a drastic transformation of representation. The basic mechanisms underlying this transformation exist in many brain areas and across phyla, suggesting that many of the principles described here could be of general relevance.

Contents

Acknowledgements	iv
Abstract	v
Table of Contents	vii
List of Figures	x
1 Introduction	1
1.1 Neural Coding	1
1.1.1 Distributed population coding	1
1.1.2 Temporal coding	6
1.2 Insect Olfactory System	9
1.2.1 Olfactory receptor neurons	11
1.2.2 Antennal lobe	13
1.2.3 Mushroom body	17
1.2.4 Synchrony and olfactory coding	18
1.3 Outline and Specific Aims	20
2 Using Noise Signature to Optimize Spike-Sorting and to Assess Neuronal Classification Quality	23
2.1 Methods	25
2.1.1 Data collection and representation	25
2.1.2 Noise model	30
2.1.3 Noise model-based clustering	33
2.1.4 Model verification tests	37
2.1.5 Sampling jitter cancellation	45
2.2 Results	48
2.2.1 Data properties	48

2.2.2	Application to real data	52
2.3	Discussion	54
2.4	Acknowledgments	57
3	Oscillations and the Sparsening of Odor Representations in the Mushroom Body	58
3.1	Results	59
3.1.1	Olfactory circuits	59
3.1.2	Resting activity	60
3.1.3	Response selectivity	62
3.1.4	Sparseness of odor representations across PNs and KCs	68
3.1.5	Mechanisms underlying sparsening	69
3.1.6	Influence of feed-forward inhibition on KC responses to odors	77
3.2	Discussion	79
3.3	Methods	82
3.3.1	Preparation and stimuli	82
3.3.2	Tetrodes	83
3.3.3	Extracellular data analysis	84
3.3.4	Responses	85
3.3.5	Sparseness	87
3.3.6	Sharp pipette recordings and staining	90
3.3.7	Immunocytochemistry	91
3.3.8	Patch-clamp recordings	92
3.3.9	Picrotoxin injections	92
3.4	Acknowledgments	93
4	Projection Neuron Population Activity: Dynamics and Coding	94
4.1	Results	95
4.1.1	Odor evoked dynamics	95
4.1.2	Spatio-temporal patterns as trajectories in PN space	101
4.1.3	Information content of single trials	107
4.1.4	Local field potentials	110
4.1.5	Kenyon cell responses	114
4.2	Discussion	116
4.2.1	PN population responses as trajectories in phase space	116
4.2.2	Comparison to the winner-less competition framework	119
4.2.3	How should one decode the PN trajectories?	120
4.2.4	When is decoding of PN activity vectors most informative?	121
4.2.5	Comparison to Kenyon cell response profiles	123

4.3	Methods	125
4.3.1	Preparation and stimuli	125
4.3.2	Electrophysiology	125
4.3.3	Data analysis	127
4.3.4	Local field potential and spike phase	128
5	Concluding Remarks	129
5.1	Spike-Sorting with Quality Control	129
5.1.1	Summary of results	129
5.1.2	Significance of results	130
5.1.3	Future directions	131
5.2	Population Coding in the Locust Olfactory System	132
5.2.1	Summary of results	132
5.2.2	Significance of results	134
5.2.3	Open questions	136
A	Projection Neuron Baseline Statistics	138
A.1	Results	138
A.1.1	Measuring firing rate statistics	138
A.1.2	Modeling PN baseline firing as a renewal process	143
A.1.3	Measuring PN-PN correlations	149
A.2	Discussion	152
A.2.1	Modeling PN firing statistics	152
A.2.2	PN-PN correlations	154
A.3	Acknowledgements	155
	References	156

List of Figures

1.1	Hypothetical tuning curves	3
1.2	Locust olfactory circuitry	14
2.1	Event detection and noise analysis on raw tetrode data	27
2.2	Illustration of the SD test on simulated data	40
2.3	Illustration of the χ^2 test on simulated data	42
2.4	Sampling jitter	46
2.5	Test of second order noise model	51
2.6	Example of classification on real data	53
3.1	Locust olfactory circuits	61
3.2	PN and KC baseline firing	62
3.3	Tetrode recordings of odor responses in PNs and KCs	64
3.4	Statistics and sparseness of PN and KC odor responses	66
3.5	In vivo sharp-electrode intracellular records from different KCs	70
3.6	Feed-forward inhibition of KCs by LHIs	72
3.7	KC responses to electrical stimulation of PNs	75
3.8	Influence of feed-forward inhibition on KC responses	78
3.9	Extracellular tetrode recordings and spike-sorting	86
3.10	Population responses and sparseness across PNs and KCs	89
4.1	PNs respond to prolonged odor presentations with constant firing	96
4.2	PN responses consist of evolving ensembles of responsive PNs	98
4.3	Visualization of PN population odor responses	102
4.4	Quantification of inter-trajectory distances	106
4.5	Odor information is available from short time slices of single trials	109
4.6	Phase-locked PN spikes show an increased odor response	111
4.7	Idealized odor trajectories	118
A.1	Example of ISI distribution and hazard function	140
A.2	Two statistical tests for PN spike trains	142
A.3	Schematic of hazard model parameters	144
A.4	Statistics from seven simultaneously recorded PNs	146
A.5	Average hazard function statistics	147
A.6	Comparison of raw and simulated PN spike trains	148
A.7	Baseline correlations between simultaneously recorded PNs	149

CHAPTER 1

Introduction

THE BRAIN is an information processing device. It takes in information about the world via an array of sensors, stores and processes the information, and then sends out signals (information) that control the movement of the body and, ultimately, behavior. Therefore, to understand how the brain works, it is critical to understand how information is represented and processed by neural circuits.

1.1 Neural Coding

1.1.1 *Distributed population coding*

The activity of individual neurons can represent information. Early studies in electrophysiology revealed that in sensory systems, single neurons encode information about an external stimulus through their level of activity (Adrian, 1926; Hubel and Wiesel, 1962; Parker and Newsome, 1998). More recently, studies have shown that the activity of single neurons in motor cortex can control a motor output (Brecht et al., 2004). In insects and other invertebrates, a small number of identifiable neu-

rons are responsible for encoding and triggering an escape response (e.g., Roeder, 1948).

Nevertheless, most brain areas are made up of large populations of neurons and must represent and process more information than any single neuron can handle. For example, primate primary visual cortex encodes all the basic features of our entire visual field (color, brightness, orientation, location), and contains over 10^8 neurons. Each neuron's activity encodes the information contained in only a tiny fraction of the entire visual scene, so knowing the activity of just one neuron is of limited value. For this reason, it makes sense to consider the responses of populations of neurons as potential units of information representation.

TUNING CURVES

Two populations of sensory neurons can represent the same sensory information in very different ways. In order to describe a particular population coding scheme, it is useful to describe the way each member of the population responds under all conditions. By presenting a wide range of stimuli and recording the strength of a single neuron's response, an experimenter can build a tuning curve for that cell. For example, figure 1.1A shows the tuning curve of a hypothetical cell in response to a range of stimuli that vary along one dimension. In practice this dimension (represented by parameter p), could correspond to color, sound frequency, temperature, or any other sensory feature to which neurons will respond. In a typical neural population, different cells will have a spectrum of different (and possibly overlapping) tuning curves (figure 1.1B).

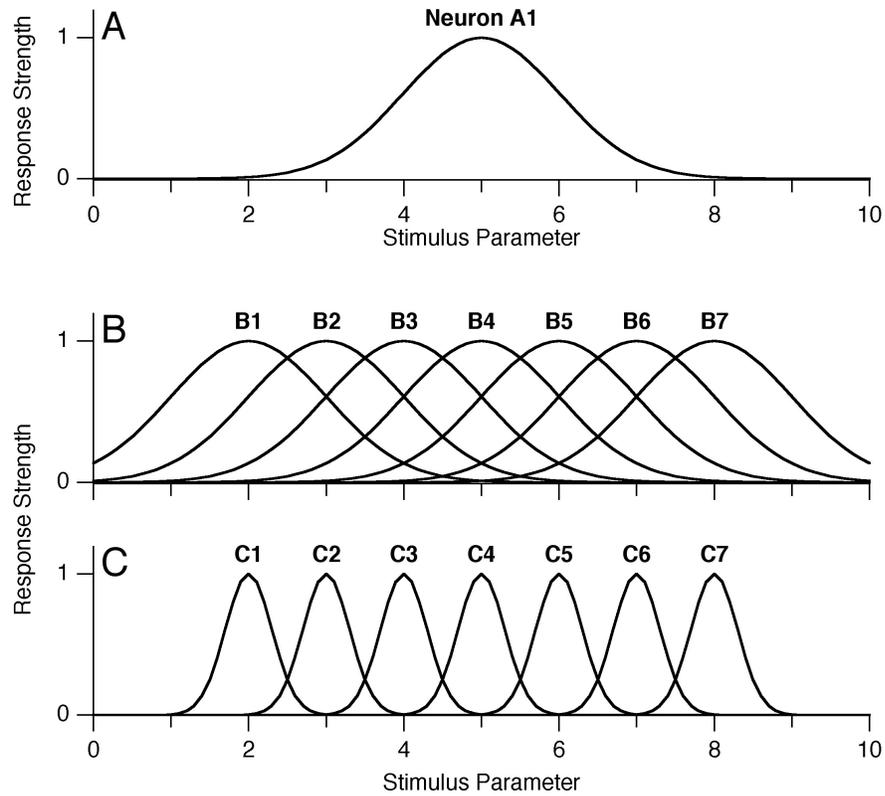


Figure 1.1. *A*, the hypothetical tuning curve for neuron *A1*. This curve measures the response intensity (which ranges from 0 to 1) for this neuron, in response to a set of stimuli, where parameter p varies from 0 to 10. *B*, the tuning curves for a population of neurons (*B1–B7*) with tuning curves similar to *A1*, but centered on different values of parameter p . *C*, the tuning curves for a population of neurons (*C1–C7*) with narrower tuning curves than population *B*.

Neural populations can be compared on the basis of their average tuning curve width. Roughly speaking, this corresponds to the range of stimuli that can elicit a response from one neuron. By comparing the set of tuning curves in figures 1.1B and 1.1C, one can see that the tuning widths of the cells in population C are narrower than those of population B. If a stimulus with parameter $p = 4$ is presented to population C, only one of the seven cells will be active (C3). In contrast, the same stimulus presented to population B would elicit some response from over half the cells. Those responses are both valid representations the stimulus $p = 4$ for those two particular populations.

In real neuronal populations as well, tuning widths vary substantially across different cell populations (Ringach et al., 1997) or even over time within the same populations (Spitzer et al., 1988). In many cases, however, it is experimentally difficult to measure tuning curves. For example, consider a neuron in the visual system that responds to images of 3-dimensional objects; it would be impossible to present all possible visual stimuli to that cell, or even a reasonable subset. In such situations, it is often more practical to measure a related statistic, population sparseness. The sparseness of a population response refers to the total fraction of cells that are active in response to a stimulus. When this fraction is low, sparseness is high. Unlike a tuning curve, sparseness can be measured for even a small set of stimuli. Returning to the simple model in figure 1.1, one can see that the response of population C is much sparser than that of B (for the reasons explained earlier). The idea of an extremely sparse neural code (e.g., a “grandmother cell” code) has

been around for decades (Barlow, 1969). More recently, experimental results have shown that sparse codes are found in areas as varied as mammalian visual cortex (Vinje and Gallant, 2000) and frog olfactory cortex (Duchamp-Viret et al., 1996).

The relative theoretical advantages of population coding with wide or narrow tuning widths (or sparse or non-sparse codes) have been examined recently (e.g., Fitzpatrick et al., 1997), but the results seem to depend critically on the specifics of the system (e.g., the noise correlation between cells (Pouget et al., 1999)). It is likely that the manner in which other neurons read out this information will strongly affect the optimal width of response tuning.

The insect olfactory system provides an ideal system for addressing many issues relevant to population codes. In the following few chapters, we will study two connected populations of neurons, both in the olfactory pathway, that employ two very different strategies of population coding. Chapters 3 and 4 will demonstrate that odor representations in the antennal lobe are distributed across a large percentage of its neurons. In contrast, in chapter 3 we show that the mushroom body, which receives direct olfactory input from the antennal lobe, has representations that are significantly more sparse than those in the antennal lobe. By closely examining the transformation that takes place between these two areas, we are able to uncover some of the mechanisms that allow a dense code to be read out and converted into a sparse code. The insect olfactory system is therefore an attractive system for future work exploring the relative benefits of sparse and non-sparse codes.

1.1.2 *Temporal coding*

How do nervous systems read out the activity of a single neuron? Up to this point we have not explicitly defined how to measure the response of a single neuron. A straightforward way of reading out a neuron's response is by computing its mean firing rate within a time window of arbitrary (or reasoned) length (e.g., a few hundred milliseconds), and tracking how this value changes across successive time windows. This measure, often described as a *rate code*, is used in a large number of experimental systems (Parker and Newsome, 1998). In sensory systems in particular, the mean firing rate of certain neurons is strongly correlated with one or more parameters of the stimulus being applied (Adrian, 1926; Hubel and Wiesel, 1962).

In a pure rate code, the only relevant measure is a neuron's mean rate of spike generation. The precise timing of any one action potential is not considered important.¹ Nevertheless, there is an increasingly large body of evidence that suggests that information can be encoded in the temporal precision of output spikes (Singer, 1999; VanRullen et al., 2005).

The term *temporal coding* is used to describe the idea that the precise temporal patterning of spiking activity may play an important role in neural information processing. Although the term is often used specifically in the context of synchrony, for the present manuscript we borrow a more general definition of temporal coding² from Dayan and Abbott (2001):

¹What *precise* really means must be defined on a case by case basis.

²A more rigorous and quantitative definition of temporal coding and an exploration of closely related terms and concepts is provided by Theunissen and Miller (1995)

Temporal coding refers to (or should refer to) temporal precision in the response that does not arise solely from the dynamics of the stimulus, but nevertheless relates to the properties of the stimulus. (Dayan and Abbott, 2001, p. 38)

Several distinct temporal coding strategies have been proposed, each with significant evidence to support its existence in areas of the nervous system.

1. SYNCHRONY

Theoretical studies have proposed that synchrony in spike timing across ensembles of neurons is used as a signal in the brain (Singer, 1999; Singer and Gray, 1995). Evidence of synchronous neural activity has been found across many brain areas in vertebrates and invertebrates alike (Adrian, 1942; Gelperin and Tank, 1990; Laurent and Naraghi, 1994; Fries et al., 2000). Nevertheless, there is still considerable debate regarding the ability of synchronous spikes to encode relevant information and about whether that information is read out by, and relevant to, other neurons.

Direct experimental evidence is difficult to collect, although some studies do address these questions. For example, Schnitzer and Meister (2003) found that more information can be extracted from synchronous spikes across retinal ganglion cells than from the spikes of single cells alone. Stopfer et al. (1997) demonstrated that disrupting synchrony in the insect olfactory system leads to behavioral deficits. The same manipulation also decreases the response specificity of the desynchronized cells' targets (MacLeod et al., 1998; Perez-Orive, 2004).

2. SYNFIRES CHAINS

Abeles (1991) proposed the model of a *synfire chain* that can generate repeatable patterns of spatio-temporal activity in response to an initial stimulus. In this model, a neural population is wired up so that a small ensemble of synchronously active neurons at one time will, after a short delay, consistently elicit the synchronous activity of a different ensemble of neurons. Such a network can sustain self-perpetuating sequences of activity that are repeatable for the same initial conditions. Evidence for such self-perpetuating spatio-temporal patterns of activity has been found in cortex (Abeles et al., 1993; Ikegaya et al., 2004), though the model remains controversial.

3. FIRST-SPIKE TIMES

Yet another type of temporal code considers the order in which spikes are fired across a population of cells. Theoretical (Thorpe, 1990, 2001) and experimental (Johansson and Birznieks, 2004; Petersen et al., 2001) evidence suggest that a sensory stimulus can be decoded by ranking a population of neurons by the order in which their first stimulus-evoked spike was fired. A first-spike based mechanism should decode a stimulus faster than a firing rate based model.

One theme common to all these temporal coding mechanisms is that spike times are measured relative to the timing of spikes from other neurons in the population. The absolute timing of a spike with respect to the stimulus is not directly important. This is a critical feature for a biologically plausible coding scheme because neurons in the brain only have direct access to the spike times of other neurons, not to

the timing of the stimulus itself. This is also an important consideration when designing an experiment to study temporal coding. Only by recording multiple neurons simultaneously can one be sure of the relative timing between their spikes.

Because spike timing is relative, temporal coding often requires a population response.³ Temporal information is only relevant with respect to the rest of the population. Conversely, by using temporal coding, the total information encoded by a population of neurons can potentially be more than the information that could be decoded from each neuron individually.

The work described in the body of this dissertation will explore a system that uses both population and temporal coding—the insect olfactory system. The next section of the introduction will provide a brief overview of this system.

1.2 Insect Olfactory System

As a model for studying neural coding, the insect olfactory system offers many benefits. One primary reason for studying insect nervous systems is their relative simplicity. Although insects can engage in many complex behaviors, including associative learning (Heisenberg, 1989; Menzel and Muller, 1996), navigation (Wehner and Menzel, 1990), and even understanding some high-level concepts (Giurfa et al., 2001), their nervous systems are much smaller and simpler than mammalian systems. For example, the first stage of olfactory processing in insects, the antennal

³Information can also be encoded in the relative timing among spikes from the same neurons (e.g., Rieke et al., 1997).

lobe, consists of only two main cell types,⁴ totalling ~1000 neurons in locust (Leitch and Laurent, 1996). In contrast, the analogous structure in mammals, the olfactory bulb, contains twice as many major neuron types and well over one million total neurons in rabbit⁵ (Shepherd et al., 2004). The large size of many insect neurons, as well as their ease of access, allows an experimenter to make *in vivo* recordings from many different types of neurons and even record from large numbers of neurons simultaneously.

The olfactory system of insects is an especially attractive system because of its strong functional and structural homology to the mammalian olfactory system (for details, see Hildebrand and Shepherd, 1997). In both systems, receptor neuron axons exhibit similarly precise and convergent targeting to specific locations (glomeruli) in the first olfactory relay (Mombaerts et al., 1996; Vosshall et al., 2000). This relay, the antennal lobe (AL, insects) or olfactory bulb (OB, mammals), is a site of significant convergence and divergence in both systems. The AL and OB contain at least an order of magnitude fewer neurons than both the number of receptor neurons that provide their input, as well as the number of cells in their main output regions (mushroom body, piriform cortex) (Hildebrand and Shepherd, 1997). In addition, both insects and mammals (as well as species from other phyla) exhibit odor-evoked oscillatory activity (Laurent and Naraghi, 1994; Adrian, 1942; Gelperin, 1999). These similarities provide reason to believe that the coding princi-

⁴In moths and some other species, the projection neurons are further classified into a number of subtypes (Homberg et al., 1989).

⁵The rabbit has 150,000 mitral and tufted cells (Allison, 1953) and the number of intrinsic cells may be at least an order of magnitude more (Shepherd, 1972; Shepherd et al., 2004)

ples revealed in insects will not only improve our understanding of neural coding in general, but may apply directly to our understanding of olfaction in mammals.

1.2.1 *Olfactory receptor neurons*

RECEPTOR MORPHOLOGY

Olfactory stimuli are transduced by a class of cells known as olfactory receptor neurons (ORNs). In insects, ORNs are located in chemosensory sensilla, small cuticular structures, which are found along the antenna, on the maxillary and labial palps, and elsewhere as well. In locusts, however, olfactory sensilla are found primarily on the antenna (Chapman, 1998).

Antennal sensilla can be classified into at least four morphological categories, although not all contain ORNs. In locust, basiconic and trichoid sensilla are known to be olfactory, and there is evidence that some types of coelconic sensilla also contain chemosensory neurons (Ochieng et al., 1998; Chapman, 1998).

An olfactory sensillum typically consists of a hair-like or peg-like protrusion of the cuticle. The cell bodies of the ORNs are located at the base of the sensillum. A single sensillum will typically contain two or more ORNs. In locust, some basiconic sensilla are innervated by up to 50 ORNs (Ochieng et al., 1998). The ORNs send dendrites towards the distal end of the sensillum. The surface of the sensillum contains a number of pores, 10–25 nm in diameter, that allow odorant molecules to enter the sensillum and contact receptors on the the outer segment of the ORN dendrites⁶ (Chapman, 1998).

⁶Odorant molecules must first pass through the receptor lymph before reaching the receptor

The axons of ORNs project down the antenna and innervate the antennal lobe. In locust these projections are restricted to the ipsilateral antennal lobe, where they may terminate in a small number of compartments, known as glomeruli (see section 1.2.2), within the AL (Hansson and Anton, 2000; Hansson et al., 1996).

RECEPTOR PHYSIOLOGY

The electrophysiological responses of individual ORNs are typically recorded extracellularly, by inserting an extracellular electrode into the base of a sensillum (Kaissling, 1995). During such a recording, the action potentials from multiple ORNs can be detected, and these are often easily discriminable on the basis of their shape (de Bruyne et al., 2001).

When ORNs respond to a square pulse of odor, they do so by reaching a steady tonic firing rate. In many cases, this constant rate is preceded by a transient increase in firing above the eventual tonic rate (Kaissling, 1971). In addition, some inhibitory responses have been observed (Kaissling, 1971; de Bruyne et al., 2001), although less is known about these responses. Unfortunately, very little data currently exist on ORN responses in the locust (see Hansson et al., 1996, for one example).

The responses of ORNs are highly odor dependent. Most insect ORNs tend to respond to a small, specific subset of odors (Chapman, 1998; de Bruyne et al., 2001). In some cases, ORNs respond almost exclusively to one behaviorally relevant compound, like a pheromone, or the odor of a host plant (Kaissling, 1971).

The specificity of an ORN's response is conferred, at least in part, by the odor-dendrites.

ant receptor (OR) proteins that it expresses. Recent genetic studies in fruit flies described a family of genes that encode all olfactory receptor proteins (Clyne et al., 1999; Vosshall et al., 1999). A single ORN expresses only one of these ORs in its dendritic membrane⁷ plus a second ubiquitous OR, present in most ORNs (Dobritsa et al., 2003; Larsson et al., 2004).

1.2.2 *Antennal lobe*

INPUTS

The antennal lobe (AL) is the first olfactory relay in the insect brain. Like its mammalian counterpart, the olfactory bulb, the AL is composed primarily of spherical regions of neuropil called glomeruli. The number of glomeruli per AL vary across different orders of insects. The locust AL contains ~1000 glomeruli (Ernst et al., 1977). In comparison, fruit flies have only ~43 (Laissue et al., 1999).

As described in section 1.2.1, glomeruli are the synaptic targets of ORN axons. The input to each glomerulus is highly specific: in fruit flies, all ORNs expressing the same OR gene have been shown to target the same one or two glomeruli in the AL (Vosshall et al., 2000; Mombaerts et al., 1996).

LOCAL NEURONS

Within the glomeruli, ORNs synapse onto the two main cell types of the AL, local neurons (LNs) and projection neurons (PNs) (Ernst et al., 1977). Local neurons have neurites that are completely restricted to within the AL. Several different classes

⁷One recent study has found an exception to this rule, a class of ORNs where at least one ORN type expresses two functional OR genes (Goldman et al., 2005).

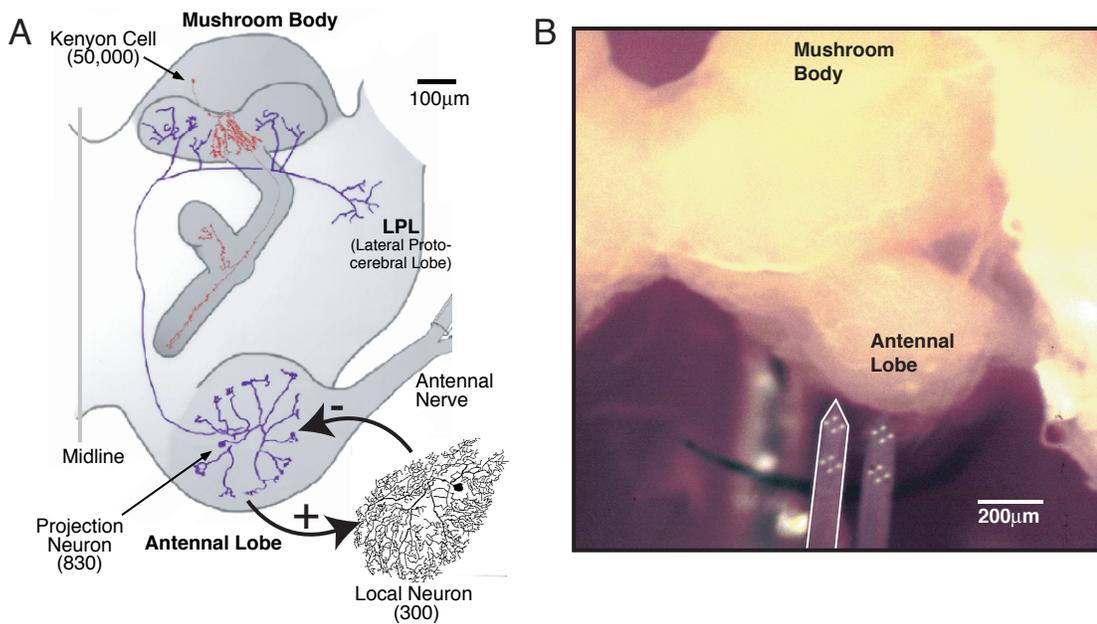


Figure 1.2. *A*, schematic of olfactory cells in the locust brain. The antennal lobe contains 830 excitatory projection neurons (PNs) and 300 inhibitory local neurons (LNs). The mushroom body contains approximately 50,000 Kenyon Cells (KCs). [Figure adapted from Wehr et al. (1999); MacLeod and Laurent (1996).] *B*, silicon multi-electrode probe alongside a locust brain. This probe is used to make extracellular recordings from PNs in the antennal lobe. Because LNs are non-spiking cells, all recorded spikes must come from PNs. Kenyon cells (KCs) and local field potentials (LFPs) were recorded in the mushroom body using either silicon probes or wire tetrodes.

of LNs have been observed across species, some innervating the entire AL, some innervating all the glomerular regions, and others innervating only a small fraction of all glomeruli (Hansson and Anton, 2000). The locust AL contains ~300 LNs, many of which (probably all) arborize widely across the AL (figure 1.2) (MacLeod and Laurent, 1996).

Local neurons in different insects exhibit a variety of physiological characteristics. In some species, like bees and moths, LNs generate sodium action potentials. In other species, like locusts, LNs show only graded potential responses and calcium (TTX-resistant) spikes (Leitch and Laurent, 1996; Laurent and Davidowitz, 1994). In locusts, as in most other insects, LNs are GABAergic (Leitch and Laurent, 1996; Hansson and Anton, 2000), although there are some reports of excitatory LNs in other species (Homberg et al., 1989).

PROJECTION NEURONS

Projection neurons form the sole output of the AL. In locusts, PNs are cholinergic and project ipsilaterally, both to the calyx of the mushroom body and to the lateral protocerebrum (figure 1.2) (Oleskevich, 1999; Ernst et al., 1977). Projection neurons show similar projection patterns in other insects including fruit flies, cockroach and moth (Homberg et al., 1989), though not all PNs in these species project to both locations (Homberg et al., 1989). In fruit flies and bees, for example, some PNs project only to the lateral protocerebrum (Wong et al., 2002; Flanagan and Mercer, 1989).

PNs receive their synaptic inputs in the glomeruli of the AL. All locust PNs have

multiglomerular dendritic trees. Locust PN dendrites seem to be very precisely organized, typically innervating 10–14 co-planar glomeruli, all at roughly the same radius from the center of the AL (Farivar, 2005). The relationship between the precise multiglomerular targeting of locust PNs and the multiglomerular targeting of locust ORNs is not yet known.

PNs in other insect species can be both uniglomerular and multiglomerular (Hansson and Anton, 2000; Homberg et al., 1989). Populations of both uni- and multiglomerular PNs have been identified in moth (Homberg et al., 1989), bee (Fonta et al., 1993), and cockroach (Malun et al., 1993).

The electrophysiological responses of locust LNs and PNs will be described in section 1.2.4.

CENTRIFUGAL NEURONS

In some species of insects, a small number of neurons have been identified that have axon terminals in the AL and dendritic trees in other regions of the nervous system (Homberg et al., 1989; Hansson and Anton, 2000). These neurons typically show traces of biogenic amines and are therefore believed to mediate modulatory feedback (Hansson and Anton, 2000). In honeybees, an identified octopaminergic centrifugal neuron, VUMmx1, is thought to mediate the encoding of an unconditioned stimulus (sucrose solution) in an olfactory association task (Hammer, 1993). At least one octopaminergic centrifugal neuron is known to innervate the locust AL as well (Bräunig, 1991).

1.2.3 *Mushroom body*

MUSHROOM BODY MORPHOLOGY

The mushroom body (MB) is a synaptic target of the PNs. The MB is primarily composed of intrinsic cells called Kenyon cells (KCs). In locust, KC cell bodies are quite small (3–8 μm , Laurent and Naraghi, 1994), and ~50,000 KCs are tightly packed above the MB calyx, the input region of the MB (see figure 1.2). Each KC sends a primary neurite into the cup-shaped calyx. There the neurite bifurcates and forms a dendritic tree spanning a small fraction of the entire calyx (Farivar, 2005). Projection neuron axons terminate throughout the main calyx (Farivar, 2005), which suggests that the majority of KCs receive some PN input. In other insect species, KCs are known to receive calycal inputs from other modalities as well (Gronenberg, 1999; Strausfeld et al., 1998).

Kenyon cell axons leave the calyx and proceed to the pedunculus, a dense bundle of KC axons, which show conspicuous reciprocal synaptic connections between neighboring axons (Leitch and Laurent, 1996). At the end of the pedunculus, KC axons bifurcate and terminate in both the α - and β -lobes, where they form synapses onto MB extrinsic neurons. These extrinsic neurons exhibit a wide variety of branching patterns. Many of these neurons have been shown to innervate the MB calyx or pedunculus as well as the lobes (Farivar, 2005; MacLeod et al., 1998; Li and Strausfeld, 1999), and could potentially mediate feedback.

MUSHROOM BODY FUNCTION

There are very few published examples of electrophysiological recordings of KCs (Laurent and Naraghi, 1994; Erber et al., 1987). Nevertheless, the MB has been strongly implicated in olfactory learning and memory (Heisenberg, 2003). In flies, MB ablation (both genetic and chemical) leads to deficits in olfactory memory (Heisenberg et al., 1985; de Belle and Heisenberg, 1994). Similarly, many mutations of genes expressed in KCs (e.g., *rutabaga*, *amnesiac*) impair olfactory learning behavior (Heisenberg, 2003).

1.2.4 *Synchrony and olfactory coding*

Recent electrophysiological studies in the locust AL and MB have revealed several basic characteristics of their odor responses. The most prominent component to the response is an odor-evoked rise in synchronized oscillatory activity at ~20 Hz. Odor-evoked oscillations have been observed in many other species (Adrian, 1942; Gelperin, 1999) and, as in these other systems, they can be observed in the locust by recording local field potentials (LFPs) (Laurent and Naraghi, 1994). Moreover, intracellular recordings from PNs, LNs and KCs all reveal odor-evoked subthreshold oscillatory activity that is phase-locked with the LFP oscillations (Laurent and Davidowitz, 1994; MacLeod and Laurent, 1996; Laurent and Naraghi, 1994).

In PNs, there is a second component to the odor response, consisting of alternating periods of excitation (spiking) and inhibition. These epochs change at a slower time scale than the oscillations. Also, unlike the odor-evoked oscillations, these

slow responses are odor- and PN-specific. Individual spikes during these PN responses all tend to lock to the same phase of the LFP oscillations (Wehr and Laurent, 1996). Thus odor-evoked spikes across the population of PNs tend to synchronize with one another.

Kenyon cells also show odor- and KC-specific firing patterns. Like PN spikes, odor-evoked KC spikes tend to lock to a single phase of the LFP (Laurent and Naraghi, 1994).

Fast GABAergic inhibition in the AL was shown to underlie the odor-evoked oscillations (MacLeod and Laurent, 1996). In the same study, MacLeod and Laurent (1996) demonstrated that by applying picrotoxin (PCT), a GABA antagonist, to the AL, odor-evoked oscillatory synchrony could be abolished. Moreover, the slow component of the PN responses, including periods of inhibition, remained unchanged. This discovery paved the way for two different studies that examined the importance of synchrony.

MacLeod et al. (1998) studied the effects of AL synchrony on odor responses in β -lobe neurons. These neurons receive olfactory input from KCs, which in turn receive direct PN connections. The odor responses of β -lobe neurons were evaluated before and after PCT injection in the AL. After injection—and therefore after the abolition of oscillatory synchrony—responses changed substantially. β -lobe neurons began responding to new odors (their tuning curves became wider), and their responses to different odors became more similar—and therefore less informative.

In a separate study in honeybees, Stopfer et al. (1997) examined the behav-

ioral consequences of synchrony disruption. Bees were trained to discriminate between pairs of odors, using a behavioral paradigm. Bees treated with PCT injections showed a significant deficit in discriminating similar odors. Their ability to discriminate dissimilar odors remained unimpaired. Taken together these results demonstrate the critical role of oscillatory synchrony in the functioning of the olfactory pathway.

1.3 Outline and Specific Aims

The focus of this dissertation is to characterize and quantify the features of the neural population code in the insect antennal lobe. Specific emphasis will be given to temporal aspects of the code. This work relies heavily on and extends previous work on olfactory coding in the locust (e.g., Wehr and Laurent, 1996; MacLeod and Laurent, 1996; Laurent, 1999; Laurent and Naraghi, 1994). These studies elegantly worked out some of the fundamental principles of the AL's odor responses, including describing two different time scales of odor-evoked temporal dynamics.

Simultaneous recordings of many PNs were necessary to precisely characterize the temporal components of the PN population response (see section 1.1.2), as well as to establish a large database of PN responses. Thus, the majority of the data presented in the following chapters were collected with multi-unit, multi-channel extracellular recordings. Multi-unit extracellular recordings present a unique challenge because no single fool-proof method exists for unambiguously extracting the activities of all neurons recorded from a set of raw data—a technique known as

spike-sorting. Adapting a rigorous spike-sorting algorithm was a necessary prerequisite to addressing neural coding questions. Chapter 2 will present several advances in the algorithms for analyzing the raw data from multi-unit, multi-channel extracellular recordings. One major advance is the development of quantitative statistical tests to assess the output quality of the spike-sorting. These tests are modular and can be adapted to other spike-sorting algorithms. A second advance presented in chapter 2 is the use of the statistics of recording-noise to optimize the clustering stage of the spike-sorting algorithm and allow it to be automated. The noise statistics are then used again in the post-sorting quality tests. The work in chapter 2 was published as Pouzat et al. (2002).

Chapter 3 quantifies some of the key differences in the population codes of the first two olfactory relays in the insect brain, the AL and the MB. The results show that the odor code in the KCs (of the MB) is significantly more sparse than that of PNs (of the AL), even though KCs receive direct excitatory input from PNs. The second part of chapter 3 addresses the mechanism underlying this striking transformation. The results point to several different underlying components, including oscillatory dynamics, feed-forward inhibition and intrinsic properties of the KCs that all work together to bring about this change. Among the conclusions of this work is that a single oscillation cycle (~50 ms) is the relevant time scale for this transformation. The work presented in chapter 3 was published as Perez-Orive et al. (2002).

Finally, chapter 4 provides the most quantitative investigation to date of the

population response of the PNs. The work in this chapter addresses some of the questions brought up by chapter 3, including the degree of PN synchrony over the course of an odor response, the fraction of PNs that respond during a single oscillation cycle, and the degree to which the PN population response changes from one oscillation cycle to the next. The results further explore the dynamics of the PN population response and describe three separate phases to an odor response. There are distinct responses to the onset and the offset of an odor pulse, both characterized by periods of strong odor-specific dynamics. Additionally, in response to odor durations longer than ~ 2 s, the PN population reaches a state of constant activity. While still to some extent odor-specific, this period is shown to be substantially less informative about odor identity than the transient response phases. The work presented in chapter 4 is in preparation for publication.

CHAPTER 2

Using Noise Signature to Optimize Spike-Sorting and to Assess Neuronal Classification Quality

UNDERSTANDING BRAIN CODES will, as a prerequisite, likely require the simultaneous sampling of large populations of neurons. While many powerful imaging techniques have been developed (e.g., membrane voltage (Wu et al., 1994); intrinsic signal (Frostig et al., 1990); fMRI (Ogawa et al., 1992)), extracellular recording remains the only method that provides both single neuron and single action potential resolution from large and distributed samples. Multi-neuron extracellular recordings, however, are useful only if the spikes generated by different neurons can be sorted and classified correctly. Although a given neuron may generate spikes with unique extracellular signal features, making the identification issue trivial, in most cases, the electrophysiologist must, from noisy and ambiguous primary data, answer the following questions:

1. What is the waveform generated by each neuron, or *unit*, on each recording site?

2. How many units were sampled by the recording?
3. On what objective basis should an individual event, or spike, be classified as originating from one of the units sampled?
4. How should superpositions, due to the nearly simultaneous occurrence of two (or more) spikes, be resolved?
5. How likely are misclassifications, that is, how often is an event generated by neuron A classified as originating from neuron B, and vice versa?
6. How can we test and quantify objectively the reliability of our classification procedure?

The first three questions have been the focus of much investigation and several methods have been proposed (reviewed by Lewicki, 1998), such as principal component analysis (Glaser and Marks, 1968), Bayesian classification (Lewicki, 1994), clustering based on the expectation-maximization algorithm (Sahani, 1999), template matching (Millecchia and McIntyre, 1978), wavelet transform based methods (Letelier and Weber, 2000; Hulata et al., 2000) and clustering methods that use spike time information to determine cluster boundaries (e.g., Fee et al., 1996a). Question 4 has been directly addressed in two studies (Atiya, 1992; Lewicki, 1994). The reliability of some of these spike-sorting procedures has also recently been tested empirically, using simultaneous extra- and intracellular recordings (Wehr et al., 1999; Sahani, 1999; Harris et al., 2000). These later studies fail to address the main concern of the present paper: how can one evaluate, from the extracellular data

alone, the reliability of the sorting procedure? The potential causes of unreliable spike-sorting are numerous; several are described in detail by Lewicki (1998). According to Lewicki (1998, p. 74), “Many algorithms work very well in the best case, when most assumptions are valid, but can fail badly in other cases. Unfortunately, it can be difficult to tell which circumstance one is in.” The simple tests we present here are an attempt to address this dilemma.

In the body of the paper, we will provide a detailed description of our methods, as well as an illustration of their use on in vivo recordings from locust antennal lobe neurons. We begin by presenting a brief description of the experimental procedure including data collection. Next, we describe the method for generating a model of the experimental noise and for testing the accuracy of the model. We then proceed to show how that model can be used first to cluster spikes, and then to test the quality of the classification. Finally, we run the entire procedure on an example of real data.

2.1 Methods

2.1.1 Data collection and representation

PREPARATION AND RECORDINGS

All experiments were carried out on adult locusts (*Schistocerca americana*) of both sexes, taken from a crowded colony and prepared as described by Laurent and Naraghi (1994).

Neurons were recorded with silicon probes from the Center for Neural Commu-

nication Technology of the University of Michigan (Drake et al., 1988). A diagram of the probe tips with 16 recording sites is shown in figure 2.1A. The probe was connected to a custom-made impedance-matching preamplifier. The preamplifier was connected to two 4-channel differential AC amplifiers (AM model 1700 AM Systems Inc.; Carlsborg, WA). The signals were bandpass filtered between 300 and 6000 Hz and amplified 10,000 times. Data were acquired at 15 kilosamples per second using a 12 bit A/D card (Win30 D, United Electronics Inc., MA).

Data with a good signal to noise (S/N) ratio were collected relatively close to the surface (50–100 μm) of the antennal lobe (AL). Spikes recorded in the AL were attributed to the activity of projection neurons (PNs), as the AL contains only two neuron populations: the PNs, which are the output neurons and fire Na^+ action potentials and the local neurons (LNs), which are axonless and fire no Na^+ action potential (Laurent and Davidowitz, 1994). We were unable to record clear spikes with the silicon probe from the antennal nerve or its projections into the AL. Afferent axons are very small and numerous (90,000), precluding clear discrimination of single neuron signal from noise.

DATA PROCESSING

Data were analyzed offline on a PC computer (Pentium III 550 MHz, 256 MB RAM) using Igor (WaveMetrics, Lake Oswego, OR) or Scilab (a free Matlab-like software package available at: www-rocq.inria.fr/scilab). All the routines were custom developed (or are freely available on the world wide web, see below) and are available upon request.

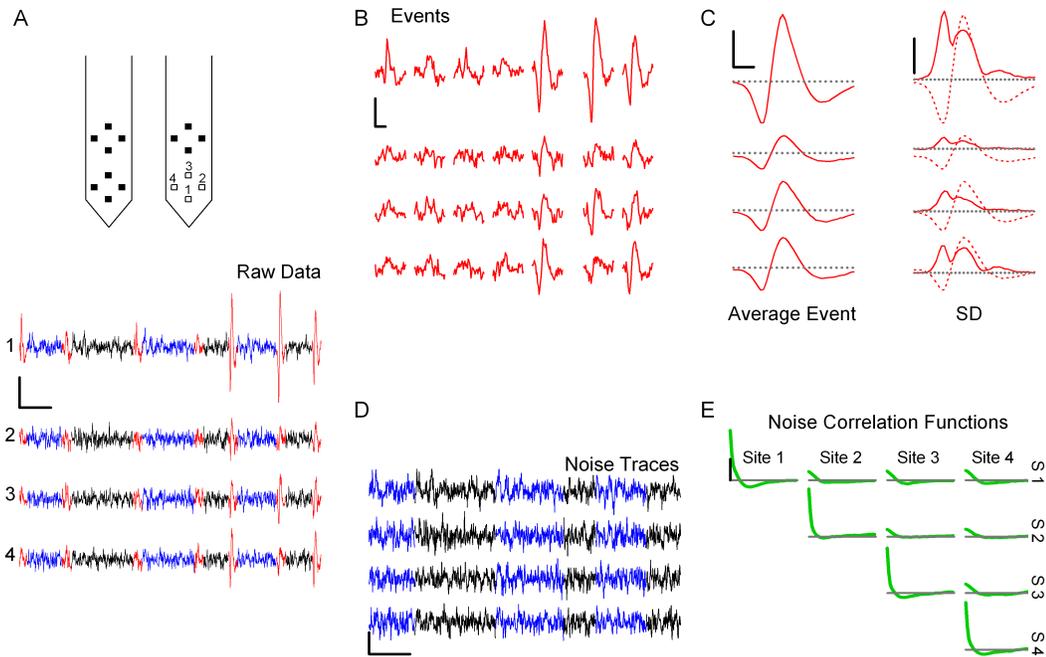


Figure 2.1. *A top*, scheme of the recording probe tips. Each shank carries two tetrodes. The four tetrodes are identical. The surface of each recording site (filled and open squares) is $177 \mu\text{m}^2$. The recording sites are placed at the corners of a square with a $50 \mu\text{m}$ diagonal. The center to center distance between two neighboring tetrodes is $150 \mu\text{m}$. The shank width is $83 \mu\text{m}$. *A bottom*, 1 s of data bandpass filtered between 300 and 6000 Hz. Site 1 is the lowest site of the tetrode; the other sites are numbered counterclockwise. Calibration: vertical, $100 \mu\text{V}$; horizontal, 100 ms. Detected events are shown in red. The traces are displayed inverted, that is, a positive deviation on the trace corresponds to a negative deviation with respect to the reference potential. *B*, the seven events shown on a smaller temporal scale (vertical, $100 \mu\text{V}$; horizontal, 1 ms). A sweep, 3 ms long, has been built with the peak of each event at 1 ms. *C left*, average event computed from the 1493 events recorded over 60 s. Vertical, $50 \mu\text{V}$; horizontal, 0.5 ms; dotted line, $0 \mu\text{V}$. *C right*, corresponding standard deviation. Vertical, $25 \mu\text{V}$; same time scale as *C left*. Dotted line: SD level expected from the noise ($15.4 \mu\text{V}$ on site 1, $15.1 \mu\text{V}$ on site 2, $14.6 \mu\text{V}$ on site 3 and $15.0 \mu\text{V}$ on site 4). Red dashed curves: average event (same as *C left*) on the same site for comparison. *D*, construction of the noise traces. After removal of 3 ms of data around each event, the remaining data are concatenated. Vertical, $50 \mu\text{V}$; horizontal, 100 ms. *E*, Noise correlation functions computed from the noise traces. Upper diagonal matrix display: auto-correlation functions on diagonal and cross-correlation functions otherwise; 3 ms sweep. Vertical, $0.01 \mu\text{V}^2$.

EVENT DETECTION

For the detection stage only, the traces were first digitally smoothed (3-point box filter). Events (i.e., putative spikes) were then detected as local maxima with a peak value exceeding a preset threshold. In cases where the spike peak occurred at slightly different times on different recording sites, only one time value was used: the time from the site with the largest peak amplitude. The detection threshold was set as a multiple of the standard deviation (SD) of the whole trace. We typically used thresholds between 2.25 and 3.5 SDs.

EVENT REPRESENTATION

Detected events can be represented in many different ways (Lewicki, 1998). Yet, the choice of a representation can strongly influence both the speed and the reliability of the classification procedure. In general, one measures a set of D parameters for each event; each event thus becomes a point in a D -dimensional space. This space is called *event space*. Our goal was to optimally predict the effect of recording noise on the distribution of points that represent events in event space. Unfortunately, several common parameter choices, such as peak and valley amplitudes or half width are computed by differentiating the raw data. This makes signal-noise separation difficult.

We therefore chose to represent each event as follows. A sweep of d consecutive sample points around the peak of the event was examined from each recording site. For our setup, we set $d = 45$ (equivalent to 3 ms), with the peak aligned to the 15th position. The sweeps were then concatenated. Therefore if one labels the

successive amplitudes of an event on site A, $A_1A_2\dots A_{45}$, on site B, $B_1B_2\dots B_{45}$, on site C, $C_1C_2\dots C_{45}$ and on site D, $D_1D_2\dots D_{45}$, the vector representing the event was:

$$\mathbf{e} = (A_1\dots A_{45}B_1\dots B_{45}C_1\dots C_{45}D_1\dots D_{45})^T,$$

where the superscript T means transpose. For our purposes, therefore, the dimensionality of the event space, D , is 180 (4×45). It will become clear that with this choice of event space, the effect of noise on the distribution of events can be easily predicted. Note that our initial peak detection for event selection introduces some sampling-induced jitter. We will ignore this for now and show later how it can be canceled.

GENERAL MODEL FOR DATA GENERATION

Following Lewicki (1994) and Sahani (1999), we use an explicit model for data generation. The general assumptions in our model are:

1. The spike waveforms generated by a given neuron are constant.
2. The signal (i.e., the events) and the noise are statistically independent.
3. The signal and noise sum linearly.
4. The noise is well described by its covariance matrix.

Assumption 1 is a working approximation, appropriate for some documented cases (Fee et al., 1996b, figure 2; Harris et al., 2000, figure 4). It also applies to our recording conditions (see results). Assumptions 2 and 3, stated for completeness,

are implicit in most already available spike-sorting methods and mean that the amplitude distribution of the recorded events can be viewed as the convolution of a pure signal with the noise distribution. We can restate our hypothesis as follows: in a noise free recording, all events generated by one unit would give rise to the *same* point in event space. In a noisy recording, however, events generated by one unit would give rise to *a cloud* of points centered on a position representing the ideal waveform of the unit. The distribution of the points should be a multivariate Gaussian whose covariance matrix would be the noise covariance matrix *regardless* of the position of the unit in event space.

2.1.2 Noise model

NOISE COVARIANCE MATRIX

To measure the statistical properties of the noise, we began by removing from the raw traces all the detected events (i.e., all the d -point sweeps) and concatenating all the inter-event traces. We call the resulting waveforms “noise traces” (see figure 2.1D). The auto-correlation function was then calculated for each recording site (diagonal, figure 2.1E), as were the cross-correlation functions between all pairs of sites (figure 2.1E). These correlations were only computed within continuous stretches of noise (i.e., the discontinuities in the noise traces due to eliminated spikes were skipped). In addition to recording noise, these cross-correlations will also account for any cross-talk between recording channels (Zhu et al., 2002).

In event space the auto- and cross-correlation functions translate into the noise

covariance matrix which was built by blocks from these functions as follows (we refer here to the four recording sites as site A, B, C and D):

$$\begin{pmatrix} AA & AB & AC & AD \\ BA & BB & BC & BD \\ CA & CB & CC & CD \\ DA & DB & DC & DD \end{pmatrix} = \Gamma,$$

where each block is a symmetric Toeplitz matrix build from the corresponding correlation function (e.g., AA is a 45×45 matrix whose first row is the noise autocorrelation function on site A, AB is a 45×45 matrix whose first row is the noise cross-correlation function between sites A and B, etc). BA is symmetrical to AB .¹

NOISE WHITENING

In order to simplify calculations and reduce the computational complexity of our algorithm, we chose to make a linear transformation on our event space (and therefore on all the detected events). The transformation matrix, U , is chosen specifically so that after transformation, the variance due to noise will be uncorrelated across dimensions (i.e., the noise covariance matrix will be the identity matrix, I). Mathe-

¹For readers unfamiliar with Toeplitz matrices, we illustrate the concept using the simple case where there are only three sample points per sweep. If the auto-correlation function on site A is the vector $(\alpha \ \beta \ \gamma)$, then AA would be:

$$AA = \begin{pmatrix} \alpha & \beta & \gamma \\ \beta & \alpha & \beta \\ \gamma & \beta & \alpha \end{pmatrix},$$

that is, $AA_{i+1,j+1} = AA_{i,j}$, for $i \geq j$.

matically, U has the property that

$$\Gamma^{-1} = U^T U, \quad (2.1)$$

where Γ is the noise covariance matrix in the original event space. A transformation matrix, U , with this property will always exist as long as the covariance matrix (Γ) is symmetric and positive definite (which it is by definition). The matrix U is obtained from Γ^{-1} with a Cholesky decomposition (Brandt, 1999, pp 479–484). A critical feature of the noise-whitened event space is that if our assumption (4) is correct (that the noise is well described by its second-order statistics), then the variance due to noise will be the same in every dimension with no correlations across dimensions (i.e., the cloud due to noise should be a hypersphere).

TEST OF NOISE MODEL

To test assumption (4), we generated a large sample of d -point long events from the noise traces. These *noise events* were taken from a portion of the noise traces different from the portion used to calculate the noise covariance matrix. Since these events should contain all noise and no signal (i.e., no spikes), these points will form a cloud around the origin in the noise-whitened event space and the distribution of these points around the origin will be fully described by the true statistics of the recording noise. We can now test if the second-order noise statistics (the covariance matrix) are sufficient in describing the actual noise distribution. We do this by computing the distribution of Mahalanobis distances (just the Euclidean

distance squared in noise-whitened space), between each noise event and the origin. In a white, Gaussian distribution, the distribution of Mahalanobis values will be a χ^2 distribution with D degrees of freedom. For our data, as we will describe in the results section, this is indeed the case.

Testing the second-order statistics is not a guarantee that the noise distribution does not have significant higher-order moments. To check for this possibility, we measured the third momentum distribution from another pool of whitened noise events. We randomly selected 500 (or more) triplets of coordinates among the 180^3 possible ones (for an event space of 180 dimensions). If we write $\mathbf{n}_i = (n_{i,1}, \dots, n_{i,180})^T$ the i th sweep of the noise sample and if, for example, $(28, 76, 129)$ is one of the triplets, the third moment for that triplet is obtained as follows (assuming a noise sample of size 2000):

$$\frac{1}{2000} \sum_{i=1}^{2000} n_{i,28} \cdot n_{i,76} \cdot n_{i,129}.$$

We will show in the results that for our data, this statistic was not significantly different from zero.

2.1.3 *Noise model-based clustering*

SPECIFIC DATA GENERATION MODEL

If our first assumption about data generation is correct (that spike waveforms are constant), the distribution of events in event space, after the linear coordinate transformation (equation 2.1), should be a set of clouds of identical shapes (hyper-

spheres), each centered on its underlying unit.² Our goal is now to determine the number of such clouds and the position of their centers in event space.

To this end, we introduce a specific data generation model (M) that extends the general data generation model by specifying the number of units, K , their waveforms and their discharge frequencies. In event space, the waveforms of the K units translate into a set of K vectors \mathbf{u}_j (joining the origin to the point representing unit j , $j \in \{1, \dots, K\}$). Our goal is to find the model that gives the best explanation of the data sample $S = \{\mathbf{e}_1, \dots, \mathbf{e}_N\}$. A common and efficient way to do this is to find the model which maximizes a posteriori the probability to observe the sample actually observed, i.e., to maximize the likelihood function (Brandt, 1999; Bishop, 1995).

The likelihood function is computed under our assumptions and in the noise-whitened coordinate system as follows. We first compute the probability (density) for unit \mathbf{u}_j to have generated event \mathbf{e}_i , $p(\mathbf{e}_i|\mathbf{u}_j)$. For that we introduce the residual vector Δ_{ij} :

$$\Delta_{ij} = \mathbf{e}_i - \mathbf{u}_j,$$

then

$$p(\mathbf{e}_i|\mathbf{u}_j) = \frac{1}{(2\pi)^{\frac{D}{2}}} \cdot \exp\left(-\frac{1}{2} \cdot \Delta_{ij}^T \Delta_{ij}\right). \quad (2.2)$$

The probability P_i for the model to have generated event \mathbf{e}_i can now be written as a weighted sum of terms like equation 2.2, one for each of the K units of the

²We should expect some outliers as well, due to nearly coincident spikes.

model:

$$P_i = \sum_{j=1}^K \pi_j \cdot p(\mathbf{e}_i | \mathbf{u}_j),$$

where π_j is the probability for unit j to occur, i.e., the ratio of the number of events from unit j to the total number of events in the sample, N . The a posteriori probability to observe the sample is then, assuming independence of the N sample elements, the product of the probabilities to observe each one of them separately:

$$\mathcal{P}(S; M) = \prod_{i=1}^N P_i.$$

The likelihood function is simply the logarithm of \mathcal{P} :

$$\mathcal{L}(S; M) = \sum_{i=1}^N \ln(P_i).$$

Several iterative algorithms exist to maximize \mathcal{L} (Redner and Walker, 1984; McLachlan and Krishnan, 1997). We used the expectation-maximization algorithm (EM algorithm, formalized by Dempster et al., 1977, and introduced in the electrophysiological literature by Ling and Tolhurst, 1983). The EM algorithm is very simple in the present context, fairly intuitive (Bishop, 1995) and its convergence to local maxima has been proven for the present model (without outliers in the sample: Boyles, 1983; Wu, 1983). Moreover, for our typical data samples, outliers do not appreciably affect the speed or accuracy of the algorithm.

The standard EM algorithm finds the best model for a given number of units. It does not provide, by itself, the actual number of units, K , present in the data

sample. Several criteria have been proposed in the statistical literature to perform this task (for an overview, see Fraley and Raftery, 1998 (especially section 2.4), and Biernacki and Govaert, 1998). Among the methods we tried, however, we found that the Bayesian Information Criterion (BIC), proposed by Schwarz (1978), worked well for our data (where most clusters are well separated in event space). The BIC penalizes an increase in the number of components by subtracting from \mathcal{L} a term proportional to $\nu \cdot \ln(N)$, where N is the number of sample events and ν is the number of model parameters. We then simply keep the model with the value of K which maximizes the BIC (Fraley and Raftery, 1998).

EVENT CLASSIFICATION

Once a model is established, we attribute each event, \mathbf{e}_i to one of the K units, by finding the j that minimizes $|\Delta_{ij}|^2$. The rationale is the following: if unit \mathbf{u}_j has indeed generated event \mathbf{e}_i , then the components of the residual vector Δ_{ij} are random numbers drawn from a multivariate Gaussian distribution and the probability to observe $|\Delta_{ij}|^2 = \Delta_{ij}^T \cdot \Delta_{ij}$ is given by a χ^2 distribution with D degrees of freedom (assuming noise whitening has been performed). By choosing the unit producing the smallest $|\Delta_{ij}|^2$ we are simply choosing the unit with the highest probability to have generated the event.³

For some events, even the smallest $|\Delta_{ij}|^2$ was very unlikely given the χ^2 distribution (e.g., in the 99th percentile). In these cases, we looked for the superposition

³Strictly speaking, we should choose the unit giving the largest product $\pi_j \cdot p(\mathbf{e}_i|\mathbf{u}_j)$, but it turned out that our units were typically far apart. Therefore a unit \mathbf{u}_l which did not generate \mathbf{e}_i gave a $|\Delta_{il}|^2$ value much too large, which is equivalent to a negligible $\pi_l \cdot p(\mathbf{e}_i|\mathbf{u}_l)$ value.

of any two units, e.g., \mathbf{u}_j and \mathbf{u}_l , which gave the smallest $|\Delta_{i,j+l}|^2$ value. To this end we tested all possible pairs of units and all relative timings of their peaks. This was easily computed for we knew the entire waveform of each unit. This approach was formalized by Atiya (1992) and an alternative method to resolve superpositions has been proposed by Lewicki (1994). If, after this step, we still did not find a small enough $|\Delta_{i,j+l}|^2$, we classified the event as an outlier.

2.1.4 *Model verification tests*

No matter how much effort is devoted to optimizing model-generation and event-classification procedures, in the end it is always possible for the results of a spike-sorting routine to be sub-optimal. In many recordings there may be pairs of neurons whose spike waveforms are close enough (with respect to the size of the noise) that their events could never be accurately distinguished. In such a case, some algorithms may lump the pair into one cluster and others might split such a pair in two. In either case, an experimenter would like to detect such a situation, and if the pair of units is really inseparable, discard the spikes from those cells or treat them as multi-unit data. Furthermore, due to the complexity of the task, even the best algorithms will occasionally generate incorrect models when given reasonable data. Again, this is a situation one would like to detect.

For this reason, we developed three tests for assessing the quality of spike-sorting results on a cluster-by-cluster basis. Since we have a quantitative model of data generation, we can use it to make detailed predictions about the properties

of our classified data. Here we illustrate the tests' principles by applying them to simulated data. In the results section we will present the same test applied to real data.

Consider the simple situation in which we record from a single site and where only two units, \mathbf{u}_1 and \mathbf{u}_2 , are active. Assume also that both units fire at low rates, so that nearly simultaneous spikes from unit \mathbf{u}_1 and \mathbf{u}_2 are rare. The original waveforms of the two units (used to generate the data) are shown in figure 2.2A. During our "recording session", we sample 500 events, superimposed in figure 2.2B (left). Each event corresponds to one of the units, to which random noise drawn from a normal distribution has been added to each of the 45 sampling points. This artificial data generation procedure is such that our model assumptions apply exactly to the sample (in this case, the noise is already white). In this sample, 300 events have been generated from unit \mathbf{u}_1 and 200 from unit \mathbf{u}_2 .

For the first two tests, we will consider two potential models of data generation. In the first case, all events of the sample are (incorrectly) classified as coming from a single unit; in the second case, the data generation model contains the two units \mathbf{u}_1 and \mathbf{u}_2 , and all events are correctly classified.

THE SD TEST

The mean event and the SD computed from all 500 events of the sample are shown on the right of figure 2.2B. Note how the SD varies, reaching maxima at times when the two waveforms (\mathbf{u}_1 and \mathbf{u}_2) differ the most. Based on our initial assumptions, we would predict that this cluster of 500 events could not all come from the same

unit. If this were the case, all the spike-to-spike variance would be due to noise, which should be constant throughout the time course of the spike.

If we now split the sample into two correctly classified sub-samples, one consisting of the 300 events generated by unit \mathbf{u}_1 and the other from the 200 events generated by unit \mathbf{u}_2 , the SD computed on the corresponding subsamples is now flat, centered on the background noise level (figures 2.2C and D). This matches precisely with what our model predicts for correctly classified clusters: all the spike-to-spike variability is due entirely to noise.

In this way, we can use this as a qualitative test of both the accuracy of the model and a proper classification of the events.⁴ After the events have been classified, the SD of every cluster can be tested. Any cluster whose SD values deviate significantly from the SD of the noise can be eliminated from further analysis (or at least scrutinized more closely). In our experience, this test is quite sensitive and can routinely detect clusters that contains multiple units, even if those units are not well-separable (see projection test).

As a final note, this test will also reliably indicate if a significant number of spikes from a small unit were not detected. This situation can arise when the peak voltage of a unit's waveform is just at the spike detection threshold. In such a case, a significant percentage of that unit's events will not be detected due to noise fluctuations. The spikes from this unit that are detected will have positive noise values near the peak, and therefore less noise variability along this portion of the waveform. This situation is therefore characterized by a dip in the SD near the peak

⁴This test was initially proposed, in a different context, by Jack et al. (1981).

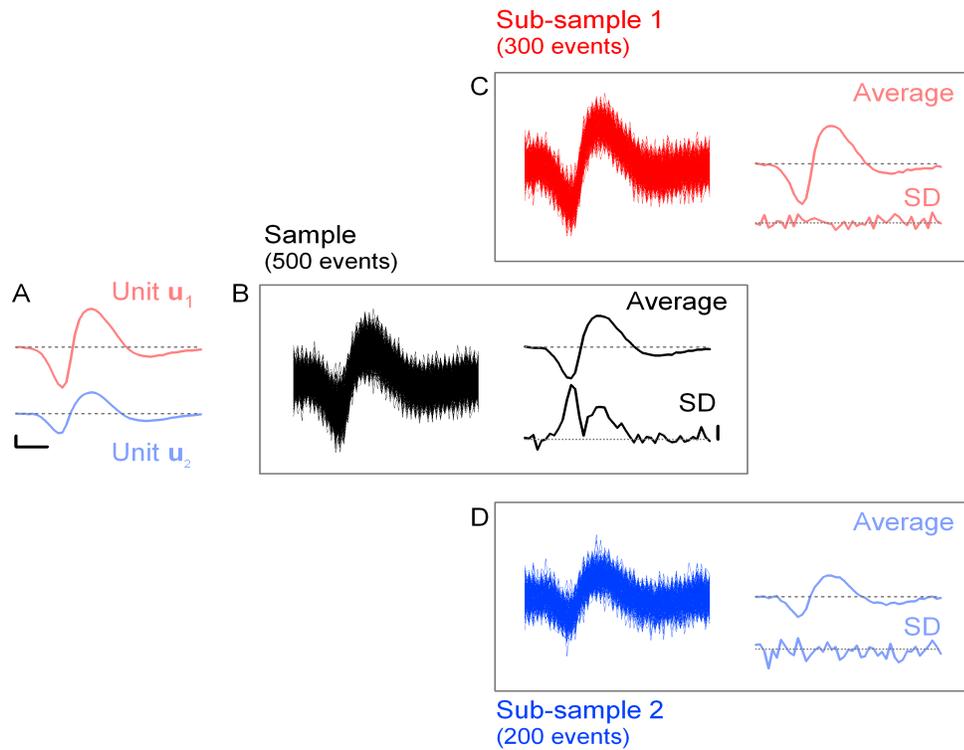


Figure 2.2. Illustration of the SD test on simulated data. *A*, waveforms of the two units used to generate the data (see text). The scale bars are arbitrary. The vertical bar has the value 1, equal to the noise's SD. To compare with real data, the length of the horizontal bar would be 0.5 ms; dashed line at zero. *B left*, 500 events generated from the two units (300 from u_1 and 200 from u_2) by adding normal white noise to the units waveforms. *B top right*, average event computed from the 500 events of the sample. *B bottom right*, SD computed from the sample. Dotted line, 1 (expectation from the noise properties). Vertical scale bar, 0.1. Notice the non-zero SD at the peak and valley of the average event. *C and D*, as in *B*, except that *C* and *D* have been built from the sub-samples generated by unit u_1 and u_2 , respectively. Note the reduction of SD variations.

of the waveform, and we routinely observe this effect empirically. Hence, a cluster that exhibits a constant SD, equal to that of the noise, is consistent with a good model together with correct spike detection and classification.

THE χ^2 TEST

In this test we test the prediction that each cluster of events forms a D -dimensional Gaussian distribution. For every unit, \mathbf{u}_j , we can compute the distance from it to all events, \mathbf{e}_i , that were attributed to it. If the prediction is accurate, the distribution of the squares of these distances should follow a χ^2 distribution with D degrees of freedom.

The test is illustrated in figure 2.3A. In the first case (one-unit model), we take the sample mean as an estimate of the ideal underlying unit. We illustrate the computation of the residual of event # 400 with such a model (figure 2.3Ai). Because we have 500 events in the sample, we obtain 500 χ^2 values. In figure 2.3Aii we plotted the cumulative distribution of these 500 χ^2 values (continuous gray curve). This empirical distribution can be compared with the expected one (dashed black curve). In this case, the expected distribution is a χ^2 distribution with $D - 1$ degrees of freedom (i.e., 44), for we have used the average computed from the same sample.

In the second case (two-unit model), we take the averages computed from the two subsamples as estimates of the underlying units (figures 2.2C and D). The classification of event # 400 is illustrated in the middle part of figure 2.3Ai. In this case, the first value suggests an unlikely event (i.e., the noise would not be

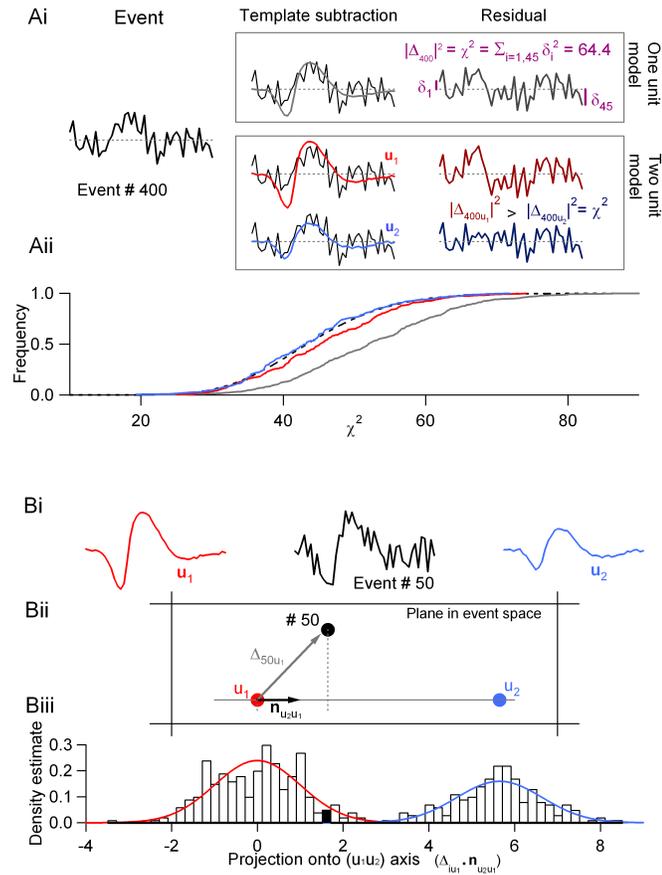


Figure 2.3. A, illustration of the χ^2 test using the same computer generated sample as in figure 2.2. *Ai*, one-unit model (top): the average event is first subtracted from event # 400, to yield the residual. The integral of the square of the residual is the χ^2 value of event # 400. Two-unit model (bottom): two units could now have generated event # 400; the waveforms of these units are given by templates u_1 and u_2 (see text). The integrals of the square of the two residuals ($|\Delta_{400u_1}|^2$ and $|\Delta_{400u_2}|^2$) are compared; the smallest indicates which one of the two units most likely generated the event, with its associated χ^2 value. *Aii*, cumulative distributions of the χ^2 values under the one- and two-unit model assumptions and their expectation (dashed line). Grey line: one-unit model ($n = 500$); red line: unit u_1 , two-unit model ($n = 300$); blue line: unit u_2 , two-unit model ($n = 200$). B, projection test. *Bi*, template u_1 (red), u_2 (blue) and event # 50 from the 500 computer generated samples. *Bii*, same objects in the plane that contains all three points. The straight line joining the two units has been drawn as well as the unit vector originating in u_1 , $\mathbf{n}_{u_2u_1}$. The vector joining u_1 to event # 50, Δ_{50u_1} (representation of the residual in event space) has been drawn as well. *Biii*, projection histogram of the 500 events of the sample. The bin containing the projection of event # 50 has been filled. The red curve is probability density function (PDF) expected from the projections of the 300 events generated by unit u_1 (60% of the sample) while the blue curve is the corresponding PDF expected from the projections of the 200 events generated by unit u_2 (40% of the sample).

expected to cause such a large deviation from the underlying unit) so the event is classified as originating from unit \mathbf{u}_2 . We thus obtain from the 500 events, two empirical χ^2 distributions (figure 2.3Aii), one corresponding to subsample 1 (red curve) and one corresponding to subsample 2 (blue curve). It is clear that these two empirical distributions are much closer to the expected one. A good classification (together with a good model) should thus yield K distributions, for a model with K units, centered on a single predictable χ^2 distribution. Like the SD test, this test is especially sensitive to the grouping of two similar-looking units into a single cluster and will produce a significant rightward shift in such situations.

THE PROJECTION TEST

According to our model assumptions, the events generated by a given unit should form a cloud of points centered on the unit in event space. The precise distribution of these points should be, after noise-whitening, a multivariate Gaussian with a covariance matrix equal to the identity matrix. Moreover, the projections of two subsamples onto the axis joining the two units which generated them should form two Gaussian distributions of $SD = 1$ centered on the two units. We can now quantitatively define the distinguishability of the two units by setting a limit on the acceptable overlap between these two distributions (overlap between distribution and event misclassifications are indeed equivalent). For instance we can decide that if more than 5% of the events coming from unit \mathbf{u}_1 or \mathbf{u}_2 are misclassified, then the two units are not distinguishable.

Figure 2.3B illustrates the procedure with simulated event # 50. As before, we

are working in a 45-dimensional space. Figure 2.3Bi shows template \mathbf{u}_1 , event # 50 and template \mathbf{u}_2 as voltage traces over time. Figure 2.3Bii illustrates the same objects in the plane in event space that contains all three. In this plane, the straight line going through u_1 and u_2 has been drawn (u is a point in event space and \mathbf{u} is a vector joining the origin to u) as well as the unit vector going from u_1 to u_2 . We can compute the projection onto the $(u_1 u_2)$ axis for each of the 500 events of the sample. Then according to our model assumptions the empirical distribution of the projections should be given by two Gaussian with $SD = 1$, centered on points u_1 and u_2 , respectively. If two units can be reliably distinguished the two distributions will not overlap. The amount of overlap between these two Gaussian is simply a function of their Euclidian distance, making it easy to convert misclassification percentage into a minimum inter-unit distance, below which a pair of units is considered nondistinguishable (e.g., assuming the same frequencies for the two units, a fraction of misclassification smaller than 5% requires a separation of 2.5 SDs between the means of the distributions; similarly, a 5 SDs separation would yield a 1% misclassification fraction).

A second feature of this test is that it can also detect whether a single unit has been incorrectly split into two different clusters. Consider the scenario where the spikes from \mathbf{u}_1 are split into two different clusters with centers (u'_1 and u''_1). The projection between these two clusters, will form a single Gaussian distribution centered at u_1 , rather than the two Gaussians predicted by the test. In this way, the projection test is most sensitive at detecting whether two clusters are inseparably

close and whether a single unit has been split among two different clusters.

2.1.5 *Sampling jitter cancellation*

One final problem to solve originates from the limited sampling frequency used during data acquisition. It is obvious that the computer's clock is not synchronized with the neurons' firing: the events will be sampled with a random delay between their peak and the nearest tick of the computer's clock. While this effect may sound like a purely theoretical concern, it can have a significant effect on the tests we describe under standard recording conditions. This sampling effect and its effects on the SD and χ^2 test are illustrated in figure 2.4. We will consider this problem in the absence of recording noise, although the problem (and its solution) exists in real, noisy recordings (see below).

An ideal waveform from a single recording site is considered in figures 2.4A and B. The ideal waveform is made of 450 points; we show two events obtained by sampling the ideal waveform once every 10 points (figure 2.4A1, only the central part of the waveform is shown). The peak of event 1 occurs at point 19 (from the origin of the sample), while the peak of event 2 occurs at point 18 (figure 2.4A1). When we build the sweeps associated with the sampled events (figure 2.1B) we align them on their peaks, causing a slight distortion, illustrated in figure 2.4A2. We see here that two sampled events arising from the same underlying waveform have different onsets and offsets. The effect of this sampling jitter on the SD is illustrated in figure 2.4A3. One hundred such events were generated by sampling

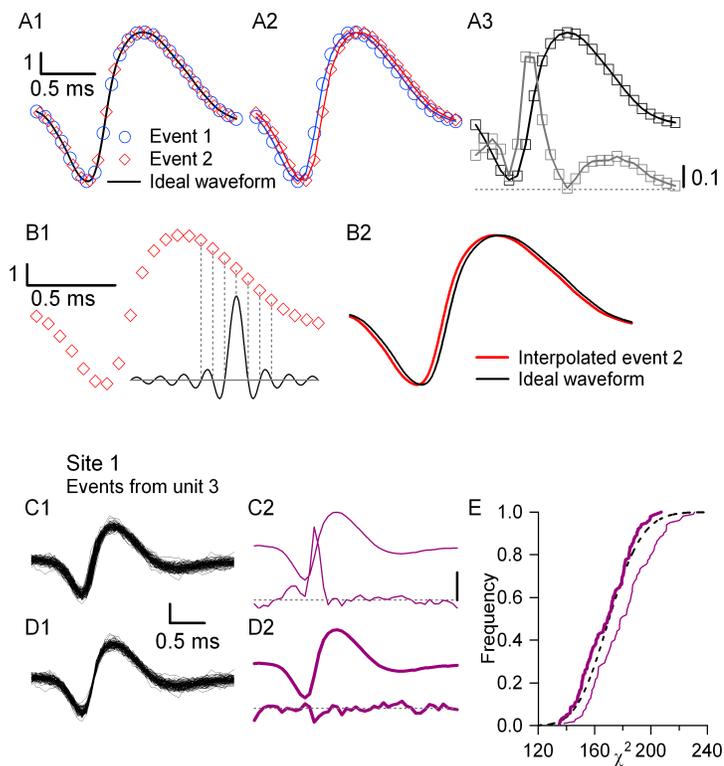


Figure 2.4. *A*, origin of the sampling jitter illustrated with simulated data. *A1*, two events are obtained by sampling an ideal waveform with two different origins. The ideal waveform is 450 points long and the two samples are 45 points long. The scale bars are arbitrary and are just labeled to help comparison with real data. *A2*, the two sampled events are aligned on their peaks and have slightly different time courses. *A3*, one hundred such events were generated from the ideal waveform (see text). Black line, mean event; grey line, SD. The dashed line is the zero SD level (no noise was added to this simulation). *B*, cancellation of the sampling jitter illustrated on event 2. *B1*, event 2 (diamonds) is exactly interpolated with a *sinc* function whose period is twice the sampling period (black trace, the peak value is 1). Continuous grey line at zero. *B2*, interpolated event 2 needs to be slightly shifted to overlap exactly with the ideal waveform. The required shift is easily obtained by minimizing the χ^2 (see methods). *C*, *D* and *E*, sampling jitter cancellation on real data. *C1*, 139 events originating from unit 3 (data in figure 2.1) on site 1, before sampling jitter cancellation. Vertical bar, 100 μV . *C2*, mean event (top) and SD (bottom). Dotted line, SD level expected from the noise properties (15.4 μV , see figures 2.1C and E). Vertical bar, 10 μV (applies only to the SD trace on *C2* and *D2*). The SD increase is slightly less pronounced than on the simulated case (*A3*) because some background noise is present and the total SD (SD_{Total} , which is the one displayed) is equal to $(\sqrt{SD_{Noise}^2 + SD_{Jitter}^2})$. *D1*, the same 139 events after sampling jitter cancellation. *D2*, mean event (top) and SD (bottom). The horizontal scale bar applies to the four graphs, *C1*, *C2*, *D1* and *D2*. *E*, χ^2 distributions before (thin line, $\langle\chi^2\rangle = 181$) and after (thick line, $\langle\chi^2\rangle = 168$) sampling jitter cancellation (performed simultaneously on the four recording sites). Dotted line, noise χ^2 distribution ($\langle\chi^2\rangle = 171$).

the ideal waveform, drawing each sweep origin from a uniform, discrete, distribution between points 1 and 10 on the ideal waveform. The sweeps were aligned on their peaks and the mean event and SD were computed. A marked increase of the SD is obvious around the times when the derivative of the mean event is significantly different from 0. It is easy to show that this SD increase is proportional to the derivative of the underlying waveform multiplied by the sampling period. This SD increase caused by the sampling jitter will also result in an increase in the χ^2 obtained after template subtraction (figure 2.3A) and corrupt our model's tests. We must therefore cancel the sampling jitter.

This is done simply by using the optimal interpolation filter to recover the full sweeps from the sampled sweeps, before realignment. This filter is the *sinc* function ($\frac{\sin(x)}{x}$) with a period equal to twice the sampling period (Papoulis, 1980). Figure 2.4B2 illustrate this interpolation procedure and its result. The 45 sample long sweep of figure 2.4B1 has been used to build a 450 points long interpolated sweep (red curve); the ideal waveform used to generate the 45 sample long sweep is shown as well (black curve) and the shift between the two curves is precisely the sampling jitter.

In practice, before running the tests we canceled jitter on every event classified as belonging to a single cluster (i.e., no outliers or superpositions). For each such event, \mathbf{e}_i , from cluster \mathbf{u}_j , we first interpolated nine points in between each true sample point to create an ideal waveform. We then aligned this ideal wave to its cluster mean by minimizing the resulting $|\Delta_{ij}|^2$. This is illustrated with real

data in figure 2.4C. Figure 2.4C1 shows the events before jitter cancellation and figure 2.4C2 shows the mean event and the SD. Figure 2.4D1 shows the same events after jitter cancellation (performed simultaneously on the four recording sites) and figure 2.4D2 the corresponding mean event and SD. The SD increase during the rising phase disappeared, replaced here by a slight over-fit (dip in the SD). Figure 2.4E shows the cumulative squared distances distributions before (thin line) and after (thick line) jitter cancellation.

Finally, caution should be used when working with short sweeps, that is, sweeps whose amplitudes at the ends are not at baseline level. Performing a *sinc* function interpolation in such cases will generate artifactual wiggles on the interpolated sweeps. If one wants to work with such short sweeps a cubic spline interpolation should be preferred.

2.2 Results

We now illustrate the generation and testing of a model using data recorded from the locust antennal lobe.

2.2.1 *Data properties*

EVENTS

A typical, 1 s long stretch of data recorded from the locust antennal lobe and band-pass filtered (see methods) is shown in figure 2.1A. Traces 1–4 originate from the four neighboring recording sites of one of the four tetrodes in the probe shown

above. Three milliseconds around each detected event have been colored red. The remainder of each trace, shown alternately in blue and black will, from now on, be considered as noise.

The seven events detected in these four traces are shown in figure 2.1B. Each sweep is 3 ms long (45 data samples) with the peak of each event at 1 ms. Figure 2.1C shows the mean (left) and corresponding SD (right) computed from all 1493 events so detected during 60 s of continuous data acquisition. The horizontal dotted line (right panel) indicates the SD expected from the background noise. Marked excesses of the SD are observed around the valleys or peaks of the mean event. Such excesses could have two nonexclusive origins:

1. Two or more units with different spike shapes are present in the data
2. A given unit generates spikes of variable amplitudes or shapes

Causes (1) and (2) should generally depend on the cell types present as well as on the region of the neuron from which the signal is recorded. Neurons producing high frequency bursts for example, often exhibit spikes of decaying amplitude during a burst (e.g., Fee et al., 1996b and Harris et al., 2000). Our model assumes that most of the SD excess can be accounted for by cause (1).

NOISE STATISTICAL PROPERTIES

The noise auto- and cross-correlation functions (figure 2.1E) were obtained as described in the methods. Note their decay time is typically less than 1 ms. This decay

time is similar to the event duration partly because the noise contains many spikes, too small to be detected as events.

VALIDATION OF THE SECOND-ORDER NOISE DESCRIPTION

It is clear that given a time dependent signal, one can always compute a correlation function (a single auto-correlation function if one records a single channel and several auto- and cross-correlation functions if one records several channels). We compute these functions in order to build the noise covariance matrix and we assume that this is a (relatively) complete description of the noise statistics. Nevertheless, one can imagine plausible scenarios where this would not be the case. If, for example, the background noise is non-stationary (e.g., Fee et al., 1996b), several noise covariance matrices could be required successively to describe the noise, while a single one could not be an accurate model. Alternatively, the noise could be stationary, but with third- or higher-order moments.

One way to test that the covariance matrix is a full description of the noise is to see how well the Mahalanobis distance distribution fits the χ^2 distribution, as described in the methods. Figures 2.5A and B illustrate the empirical cumulative distribution (2.5A) and density (2.5B) of Mahalanobis distances for an actual noise sample. The expected values for these quantities have been plotted as well, the close match between actual and expected entities suggest that the noise distribution is well approximated by a multivariate Gaussian distribution.

Although the Mahalanobis distance test is quite sensitive, it is a necessary but not sufficient test of the accuracy of our noise description. We thus performed an

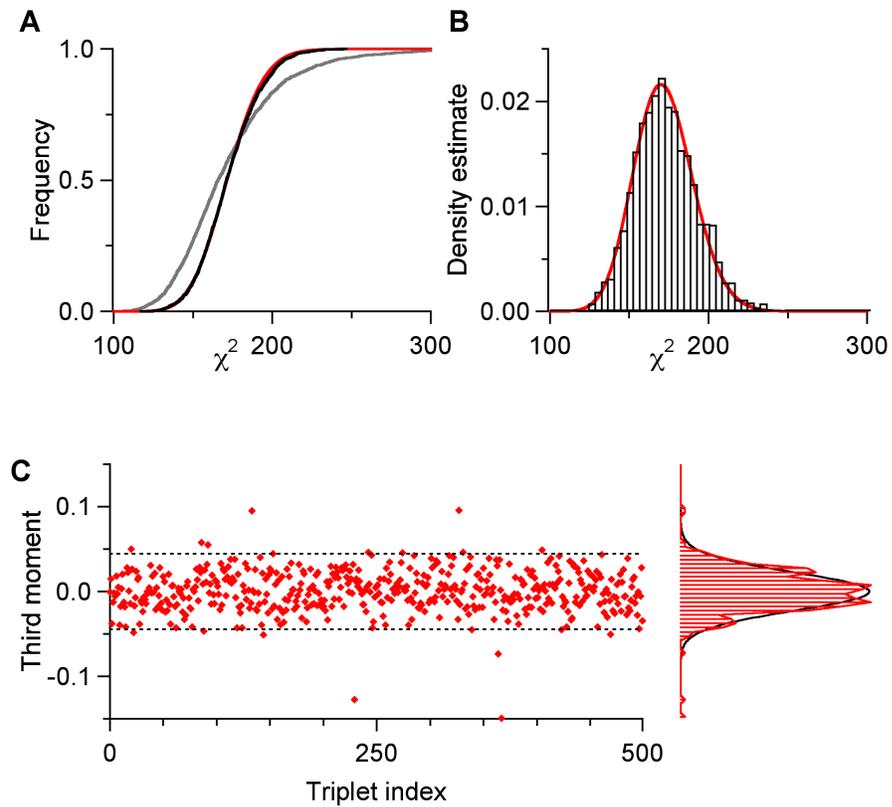


Figure 2.5. *A*, black curve, empirical Mahalanobis distance distribution obtained from a noise sample with 2000 events after noise whitening (see methods), the expected χ^2 distribution is shown in red. *B*, χ^2 PDFs. Red, expected PDF; black histogram, empirical probability density estimate from the same noise sample as in *A* and after coordinate transformation. *C*, third moment distribution of the whitened noise (see text).

additional test aimed at detecting deviations of the actual noise from its representation based on its covariance matrix. We estimated the distribution of the third moment about the mean of noise sample (see methods). Figure 2.5C shows that for 500 randomly chosen coordinate triplets the average value of the third moment (sample size = 2000) has a Gaussian distribution with 0 mean and an SD of $\frac{1}{\sqrt{2000}}$, as expected. Taken together these results suggest that a noise description based on its covariance matrix is accurate enough for our purpose.

2.2.2 *Application to real data*

The methods described can now be applied to the real data of figure 2.1. Once the specific data generation model has been obtained (see methods), it is used to classify each of the 1493 events detected during 60 s of continuous data acquisition. Of the 1493 events detected, 1361 were classified as pure events (294 events from unit 1, 391 from unit 2, 139 from unit 3, 333 from unit 4 and 204 from unit 5), 89 were classified as superpositions of two different units and 43 (i.e., less than 3%) were classified as outliers.

The pure events of three of the five units are displayed in figure 2.6B1–3 together with their mean and SD. The SD test seems to be met by the events of units 3 (figure 2.6B1) and 4 (figure 2.6B2) but not by the events of unit 5, which do not have a flat SD (figure 2.6B3). This is confirmed by the χ^2 test (figure 2.6C). The χ^2 distributions are expected to be on the left side on the noise χ^2 distribution (dotted line, figure 2.6C) for two reasons: they are computed by using the mean

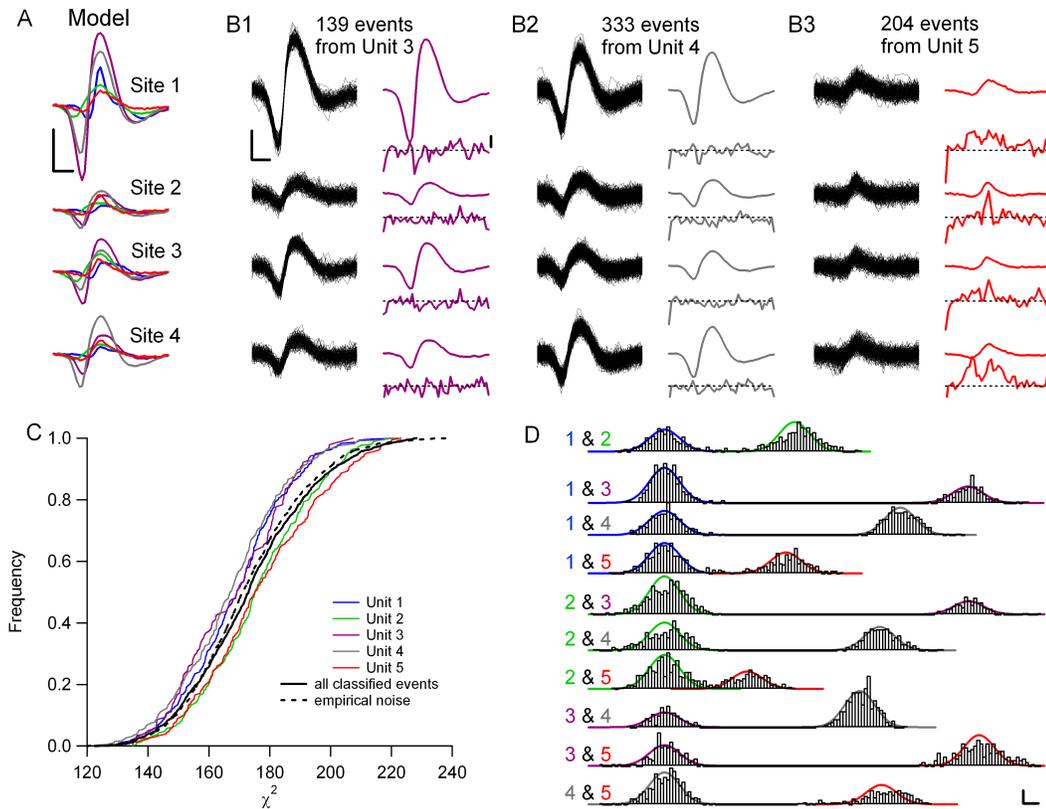


Figure 2.6. Example of classification on real data from the locust antennal lobe. *A*, waveforms of each of the five units of the model on each of the four recording sites. Vertical, 100 μV ; horizontal, 0.5 ms. *B*, events from three of the five units. *B1 left*, the 139 events generated by unit 3; *B1 right*, mean event and SD. Calibrations: left, as in *A*; right, SD of site 1, 20 μV . Dotted lines on SD graphs: expected noise (see figure 2.1C). *B2 and B3*, same as *B1*, for unit 4 and 5. All graphs in *B* have same scale. *C*, χ^2 test. Cumulative distributions of the pure events from the 5 units and from all the 1450 classified events (pure events from each of the 5 units and 89 events classified as superpositions). Dashed line, empirical noise χ^2 distribution. *D*, projection test. For each of the 10 possible pairs of units, the empirical projection of the events belonging to one or the other unit of the pair is shown (histograms). The expected probability densities are shown superposed on each graph (see text). Vertical bar, 0.1. Horizontal bar, 1. Note the absence of overlap between the projections of any two units, as well as the systematic discrepancy between the empirical and expected densities for units 2 and 5.

event of each unit as a template (i.e., 1 degree of freedom is lost) and the sampling jitter cancellation results in a slight overfit (see figure 2.4D2). On this basis, the distributions of units 2 and 5 are suspect.

The projection test (figure 2.6D) confirms the poorer quality of units 2 and 5 but shows that all units can be unambiguously distinguished. The empirical projection density estimates have been plotted on each graph (histograms) together with the expected distributions (colored curves). The projections are obtained by projecting all the pure events generated by one or the other unit of the pair. The expectations are completely defined by the knowledge of the distance between the two units of a pair and by the respective number of pure events generated by each of the two units (*e.g.*, in the top graph, the distance between the two units is 8.55, the number of events from unit 1 is 294 and the number of events from unit 2 is 391; therefore, the integral of the blue Gaussian is $294/(294+391) = 0.43$ and the integral of the green Gaussian is 0.57). Based on those tests, spikes originating from unit 1, 3 and 4 would be kept for further analysis while spikes from unit 2 and 5 would be discarded or at least taken with caution, for each of these two distributions likely contains more than one unit.

2.3 Discussion

We have shown that a very simple model can explain electrophysiological data collected by extracellular recordings in the locust AL. The combination of an accurate noise model with an explicit model for data generation leads to specific quantita-

tive tests that the classified data should meet. These tests are objective and can be graphically displayed, thus enabling the experimenter or the reader to assess the quality or trustworthiness of the analyzed data. It should be clear that the tests can be applied to the final results of any classification procedure.⁵ These tests could therefore form a basis for comparison between different spike-sorting techniques. The less-rigorous cluster cutting methods, used in particular by commercial software, sometimes leave the user or the reader with untestable confidence in the data. The adoption of objective measures such as those we propose here would, we believe, alleviate this growing problem.

Our method does not take spike timing into account at any stage. That is, no explicit refractory period is set, and no general form of the spike train autocorrelation is required, as it is in many other methods (e.g., Fee et al., 1996a; Harris et al., 2000). In this way, spike timing information can be used as another completely independent measure of quality. In our data we find very reproducible inter-spike interval distributions, among neurons in a given animal, as well as across animals (see appendix A, figure A.5), further confirming the accuracy of the procedure. Furthermore, our tests routinely detect situations which would be missed by timing-based test alone. For example, consider a cluster that contains only half of the spikes generated by a single unit, either because that unit was split between two clusters or because a substantial percentage of its spikes were not detected. While this cluster will still exhibit a normal-looking refractory period and autocorrelation, it should

⁵Two of the test require a knowledge (or at least an estimate) of the noise covariance matrix Γ , but that can always be trivially obtained from the classified events.

still be detected by the tests we propose (by either the projection test or the SD test, respectively).

The model we use here for the locust data is clearly the simplest possible one. It is not expected to hold for all data sets, however, we expect it to work successfully in a wide variety of experimental conditions. In rat neocortex, for example, Fee et al. find that for most of the units they recorded “[t]he variability of spike residuals is nearly identical with that of background activity”(Fee et al., 1996b, p. 3831, see also figures 1B, D, E, and 2 in the same article), indicating that the tests we propose should work in this system as well. Furthermore, using the framework we describe, one can introduce more sophisticated models of data generation to analyze more complicated data sets. Once a new model is specified, the statistical tests we introduced can be readily generalized. Of particular interest are spikes with nonstationary waveforms (e.g., within bursts). One way to model the waveforms of such neurons would be to assign to each unit a cluster specific covariance matrix, to account for the added variability in spike shape. These cluster specific matrices would complement a global covariance matrix describing the noise. Although the use of a covariance matrix to model spike waveform variability is only an approximation, preliminary results using in vivo data collected from the rat hippocampus (generously shared by K. Harris and G. Buzsaki) indicate that the distribution of spike waveforms from a bursty cell is well described by its second order statistics (i.e., by a covariance matrix). Another alternative would be to develop a model of the dependence of the spike waveform on the inter-spike

interval (as suggested by Fee et al., 1996b) and use this model to scale the template before computing the residual. Our tests would then be directly applicable.

Vertebrate data also often exhibit nonstationary noise (e.g., Fee et al., 1996b). Such data would require a more precise description of the noise. For instance, an extension of the current model could use a time dependent noise covariance matrix. The noise whitening would then be applied by taking into account each event's time of occurrence.

2.4 Acknowledgments

The work in this chapter is the result of a collaboration with Christophe Pouzat, who formulated the core idea. Further development of the sorting algorithm and quality tests, as well as their implementation in code, was a cooperative effort.

This chapter is reprinted from *Journal of Neuroscience Methods*, vol. 122, C. Pouzat, O. Mazor, and G. Laurent, "Using noise signature to optimize spike-sorting and to assess neuronal classification quality", pp 43–57, 2002, with permission from Elsevier.

CHAPTER 3

Oscillations and the Sparsening of Odor Representations in the Mushroom Body

ELECTROENCEPHALOGRAPH AND LOCAL FIELD POTENTIAL oscillations generally indicate periodic coherent synchronization of neuronal assemblies (Adrian, 1942; Gelperin and Tank, 1990; Gray, 1994; Bragin et al., 1995; Steriade et al., 1996; Csibra et al., 2000). While the occurrence of macroscopic oscillations has now been correlated with various sensory, behavioral or cognitive states in mammals (Gray et al., 1989; Eckhorn et al., 1988; Rodriguez et al., 1999; Patel and Balaban, 2000; Fries et al., 2000; Engel et al., 2001), the functional significance of such observations is debated (Shadlen and Movshon, 1999; Abbott and Dayan, 1999). Many hypotheses based on temporal correlations have been proposed (Abbott and Dayan, 1999; Shadlen and Movshon, 1999; Hopfield, 1995; von der Malsburg and Schneider, 1986; Diesmann et al., 1999): among others, that cortical neurons might act as coincidence detectors, rather than integrators, and thus select for correlated input (Abeles, 1982; König et al., 1996). Most hypotheses, however, remain tentative for lack of direct experi-

mental test. The olfactory nervous system, in which molecular design (Clyne et al., 1999; Gao et al., 2000; Vosshall et al., 2000; Mombaerts et al., 1996), circuit architecture (Mombaerts et al., 1996) and oscillatory dynamics (Adrian, 1942; Gelperin and Tank, 1990; Laurent and Naraghi, 1994; Lam et al., 2000) appear common across phyla, offers a rare opportunity to study some of these coding issues.

3.1 Results

3.1.1 Olfactory circuits

The insect antennal lobe (AL) is the analog of the vertebrate olfactory bulb. In locusts, each AL receives input from ~90,000 ORNs and contains ~1130 densely interconnected neurons (300 local inhibitory neurons, LNs; 830 excitatory, multi-glomerular projection neurons, PNs) (Leitch and Laurent, 1996; MacLeod and Laurent, 1996). Each AL sends distributed projections to the ipsilateral mushroom body (MB), a memory area (Heisenberg et al., 1985; Dubnau et al., 2001; McGuire et al., 2001; Zars et al., 2000). PNs are the only channel for olfactory input to the MB. Conversely, there is no evidence for feedback from the MB to the AL. Each locust MB contains ~50,000 small neurons (Kenyon cells, KCs) (Laurent and Naraghi, 1994; Kenyon, 1896), whose spiny dendrites receive direct input from PNs (Laurent and Naraghi, 1994). In locusts, each PN contacts about 600 KCs (~30 synaptic varicosities per PN axon, times ~20 distinct synaptic contacts with different KC profiles per varicosity, Leitch and Laurent, 1996). Each KC receives contacts from many PNs, as seen from incremental electrical stimulation of PNs. The total number of outputs

made by all PNs onto KCs must equal the total number of PN inputs received by all of KCs: Hence, if 830 PNs project to 50,000 KCs with 1:600 divergence, the average PN-to-KC convergence is on the order of 10. While unknown, these ratios probably vary by little more than a few fold across the PN/KC populations.¹ The dendritic tree of a typical KC contains 100 to 200 spines (Laurent and Naraghi, 1994). Many of these inputs must thus originate outside of the AL.

Odor-evoked PN responses exhibit globally coherent 20–30 Hz oscillations and stimulus- and PN-specific slow modulation of firing rate, both shaped in great part by LN-mediated inhibition (MacLeod et al., 1998; Stopfer et al., 1997; Wehr and Laurent, 1996; MacLeod and Laurent, 1996; Laurent et al., 1996). Hence, during a stimulus, the AL output consists of barrages of spikes from an evolving PN assembly. While individual PN spike timing during one oscillation cycle can be phase locked, it is not so for all PNs active during that cycle. At what time(s) a PN locks to others depends on both the odor and the PN. To understand the decoding of PN output by KCs, we examine the firing behavior of both populations at rest and in response to odors.

3.1.2 *Resting activity*

Baseline activity profiles of PNs and KCs were measured over several-minute long stretches of uninterrupted recording in naïve animals, using multiple tetrode recordings (see section 3.3.2, p. 83). At rest, the PN population fired at a mean rate

¹Recent results by Jortner, Farivar, and Laurent (submitted) suggest that the PN-to-KC convergence is closer to ~400.

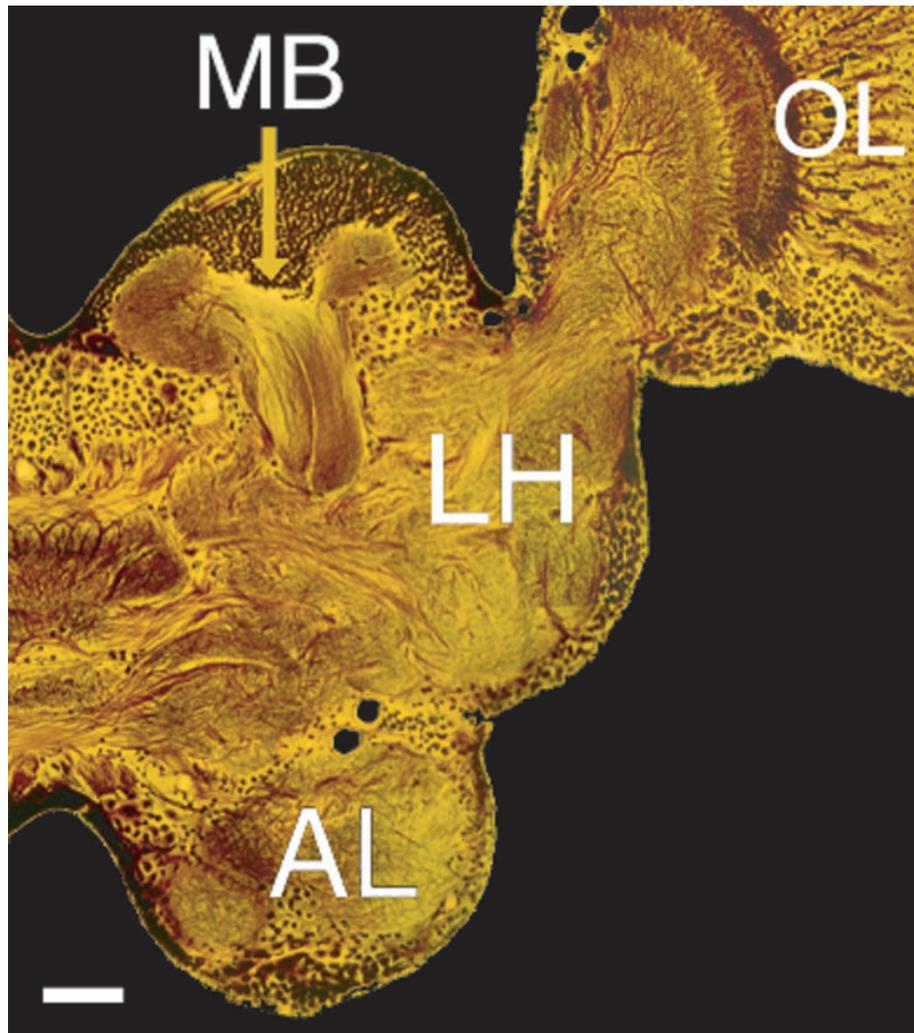


Figure 3.1. Olfactory circuits: Transverse section of the locust brain (left half, Bodian stain). Olfactory input originates from olfactory receptor neurons (ORNs) on the antenna. ORN axons terminate in the antennal lobe (AL), where projection neurons (PNs) act as relays, with projections to the mushroom body (MB) and the lateral horn (LH). OL: optic lobes. Calibration: 80 μm .

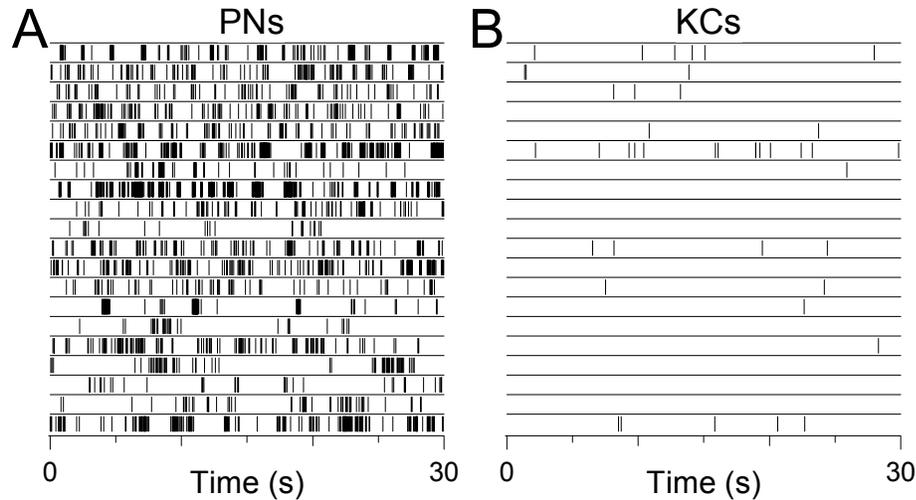


Figure 3.2. PN and KC baseline firing in the absence of odor stimulation (see text). Thirty-second rasters of 20 PNs (A) and 20 KCs (B) recorded with tetrodes. Note the exceedingly low baseline activity of KCs. (Empty rasters denote absence of action potentials during the randomly selected segment chosen for display. These rasters, however, of course originate from identified KCs, whose action potentials occurred at other times during the recording period.) [PN data (A): O. Mazor. KC data (B): J Perez-Orive.]

of 3.87 ± 2.23 spikes/s per PN (range: 0.49 to 10.4, $n = 35$ PNs). Baseline firing was >100 times lower in KCs (median: 0.025 spike/s, interquartile range: 0.088 spike/s; $n = 23$ KCs) (figure 3.2). Hence, despite a constant excitatory drive from PNs, KCs at rest remained remarkably inactive.

3.1.3 Response selectivity

PNs and KCs were challenged in awake animals with a panel of odors (typically 17; range: 5–24; 5–25 trials per odor; 1 s pulses; 20–30 s between trials; see section 3.3.1, p. 82). Experimental conditions were identical for PN and KC recordings.

SPIKING RESPONSE PROBABILITY

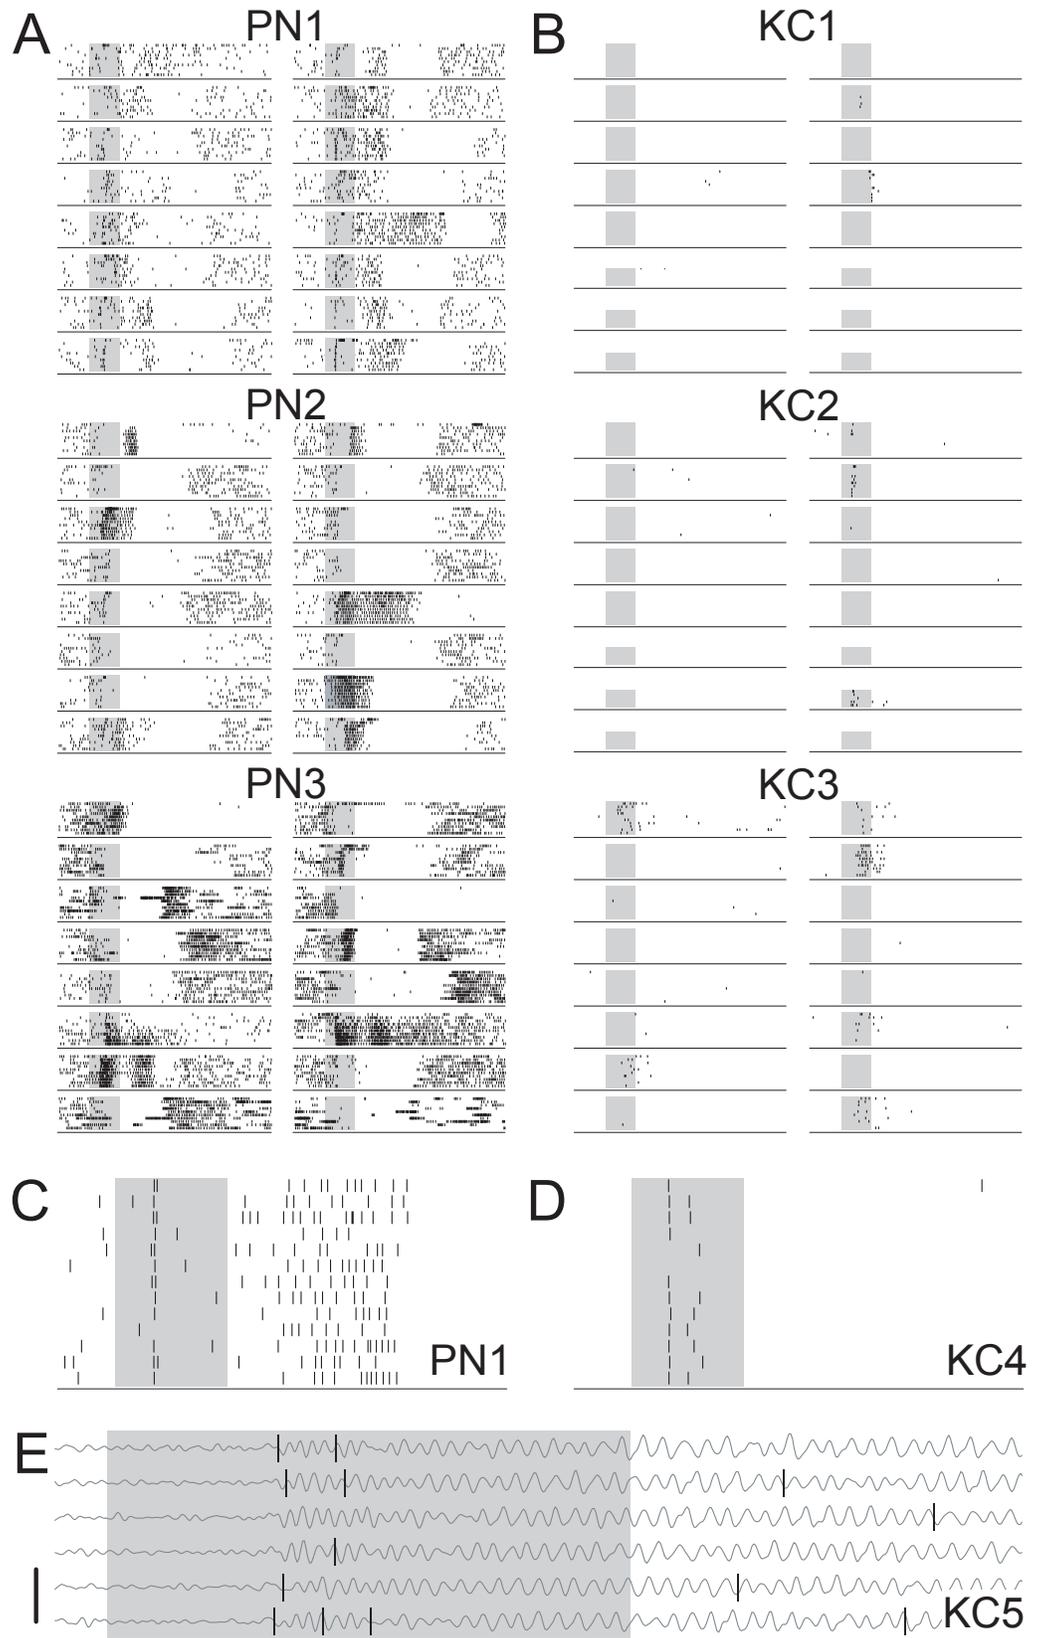
The probability of observing a stimulus-evoked change in firing behavior was different across the two populations (figure 3.3). Most PNs exhibited a reliable change in firing behavior within the first few seconds following stimulus onset. They showed complex temporal patterning (with increases and decreases in instantaneous firing rate) that often greatly outlasted the stimulus itself (figure 3.3A). Many of these responses were inhibitory, and many of these inhibitory periods were followed by a period of increased firing, up to five seconds after stimulus offset. We analyzed excitatory response probabilities across PNs (and KCs) quantitatively, using a variety of methods and analysis windows. We show here the results obtained with Method A (see section 3.3.4, p. 85). Results obtained with the other methods are nearly identical (table 3.1 and figure 3.10). The distribution of response probabilities for PNs was broad (figure 3.4A), with means over all cells of 0.64 ± 0.32 ($n = 58$ PNs, 1140 PN-odor pairs). KC responses to these same odors were extremely rare: over all KCs ($n = 74$ KCs, 1101 KC-odor pairs), 58% failed to show any detectable response to any of the odors presented (figure 3.4A). The distribution of response probabilities was heavily skewed towards low values (figure 3.4A), even when considering only those KCs that produced at least one response. The mean response probability was 0.11 after averaging all KCs' individual response probabilities (figure 3.4A) (median: 0.00; interquart. range 0.12). Figure 3.3B shows three typical responsive KCs. Among all recorded KCs, only two responded to all odors presented (10 and 12 odors, respectively). To avoid possible sampling bias, recordings were made

from all regions and depths of the KC soma layers. Responsive and unresponsive KCs were found everywhere, consistent with the anatomy of PN axonal projections in the MB (Laurent and Naraghi, 1994). Similarly, no selection bias towards strong responses existed, for the great majority of them were extremely brief (c.f., below) and rarely detected on line. Selective and promiscuous KCs could occur simultaneously on the same tetrode, indicating that differences in tuning width were not caused by global modulation of excitability over time.

RESPONSE INTENSITY

Response patterns and intensities differed in PNs and KCs. While PN responses often lasted several seconds (figure 3.3A), KC responses were brief and lacked the slow temporal patterning typical of PNs (figure 3.3B). Using responsive cell-odor pairs, we counted action potentials produced by PNs and KCs over the 3 s window after stimulus onset. The distribution of PN spike counts over that period was broad, with a mean of 19.53 ± 10.67 spikes. KCs responded with 2.32 ± 2.68

Figure 3.3 (on the next page). Tetrode recordings of odor responses in PNs (*A and C*) and KCs (*B, D, and E*). In all panels, shaded area = odor puff = 1 s. *A*, responses of three simultaneously recorded PNs (PN1–PN3) to 16 different odors (first 10 trials with each stimulus displayed). Odors from top, left column: hpo, don, che, hx3, unn, min, oca, pnn; right column: chx, oco, nnn, thx, 2hp, nna, 3hp, hxo; abbreviations in methods, p. 82). *B*, responses of three KCs to the same 16 odors. Conditions as in *A* except: for six of the odors, KC1 and KC2 have only five trials; in KC2, the 7th odor in the right column is hxa. *C*, expanded view of PN1 raster in response to hxo (trials 3 to 15). Note alignment of spikes. *D*, response of a fourth KC to hx3 (trials 3 to 15). Note low baseline activity and alignment of first spike in the response across trials. *E*, response of a fifth KC with superimposed LFP, recorded in the MB (10–55 Hz bandpass). Note phase-locking of KC spikes. LFP: 200 μ V. [PN data (*A, C*): O. Mazor and S. Cassenaer. KC data (*B, D, E*): J Perez-Orive.]



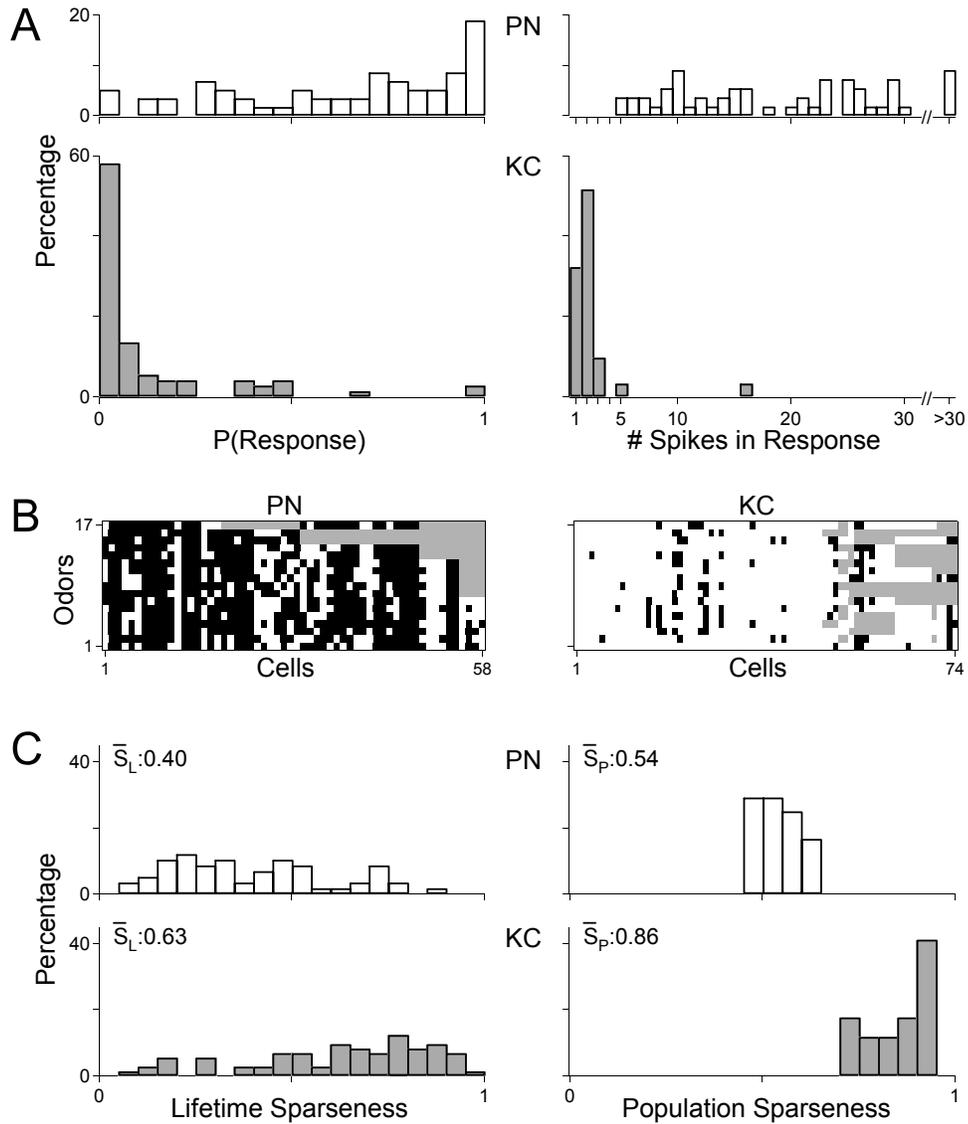


Figure 3.4. Statistics and sparseness of PN and KC odor responses (see methods). *A left*, probability of responding to $x\%$ of all odors tested (x in 5% bins) (see section 3.3.4, p. 85); note opposite skew in PN and KC distributions. *A right*, response intensity distributions (1 spike bins, measured over 3 s window). Spike counts were computed only from cell-odor pairs with a detected excitatory response during the analysis window. *B*, excitatory responses (filled squares) of individual PNs and KCs (columns) ($n = 58$ PNs, 74 KCs) to 17 different odors (rows: hx3, thx, chx, hxo, hpo, oco, nna, nnn, don, pnn, 2hp, 3hp, oca, unn, che, min, hxa; abbreviations in section 3.3.1, p. 82). Open squares denote inhibition (PNs only) or absence of a response. Grey squares: not tested. *C*, distributions of lifetime (left) and population (right) sparseness, computed across all cells and all tested odors. S_L and S_P are significantly different across PNs and KCs ($p < 0.001$, t' -test for S_P , z statistic for S_L). [PN data: O. Mazor and S. Cassenaer. KC data: J Perez-Orive.]

spikes² (figure 3.4A). We found a negative correlation between KC spike count and response selectivity (Spearman ranked correlation coefficient: -0.567 , $p < 0.05$).

TEMPORAL PRECISION

PN spike probability and precision is PN-, odor- and time-specific (MacLeod et al., 1998; Stopfer et al., 1997; Wehr and Laurent, 1996; MacLeod and Laurent, 1996; Laurent et al., 1996). Time-locked PN spikes were easily detected when they occurred in isolation (e.g., figure 3.3C), but were found also within sustained responses, consistent with previous intracellular results. In KCs, individual responses typically contained about two spikes (figure 3.4A), one of which at least could be precisely locked to stimulus onset with a fixed delay. Stimulus-locked spikes were often the first ones in the KC's response, but could occur at any cycle. The first spike in the response of KC4 (figure 3.3D), for example, had a jitter of only ± 4 ms relative to stimulus onset. Stimulus-locked spikes with such small jitter, however, were not commonly observed. Another measure of precision, more relevant to this system, is the timing of each action potential relative to its LFP oscillation cycle (phase) (figure 3.3F). The mean phase of KC spikes was $83^\circ \pm 77^\circ$ ($n = 18$ KCs; where 0° is oscillation peak, figure 3.6F). Mean spike phase was the same in the most as in the least specific KCs ($90^\circ \pm 67^\circ$ vs. $86^\circ \pm 81^\circ$, $n = 5$ cells each). The spikes within a doublet (or triplet) were always separated by one to a few oscillation cycles (e.g., figure 3.3F). This indicates that appropriate PN drive to individual KCs lasted several oscillation cycles, and that when a KC spike was fired, it occurred preferentially

²Most KC spikes occurred in the beginning of the response: response intensity was 2.33 ± 2.02 spikes over the first 1.4 s, while PNs produced 12.84 ± 7.29 spikes on average in that period.

at the same phase of its oscillation cycle.

3.1.4 Sparseness of odor representations across PNs and KCs

Figure 2B compresses the responses of 58 PNs and 74 KCs to the same 17 odors and illustrates the contrast between the two population representations.³ A simple estimate of population sparseness (S_p) is the proportion of cells unresponsive to each stimulus, averaged over all stimuli. It thus represents the sparseness of the representation of each odor across the population, averaged over all odors, but ignores the strength of each response. S_p was 0.90 in KCs and 0.33 in PNs. S_p can also be calculated more directly using the neurons' firing rate distributions (Willmore and Tolhurst, 2001) for each tested stimulus, whether we detected a response or not (see section 3.3.4, p. 85). Applied to PNs and KCs, this measure of S_p was again always greater in KCs (figure 3.4C). Finally, sparseness can be calculated for each cell across all the stimuli it has experienced. This measure, called lifetime sparseness, S_L , approximates the mean tuning width of each neuron, averaged over all neurons. Again, S_L was significantly higher in KCs than in PNs (0.63 vs. 0.40, $p < 0.001$, t' -test, figure 3.4C). S_L and S_p were also calculated using the other response analysis windows, or using only the odor responsive cells. By all measures, odor representations were always significantly sparser across KCs than PNs (table 3.1 and figure 3.10).

³Responses were determined here according to Method A, (see section 3.3.4, p. 85). Nearly identical results were obtained if responses were assessed using different criteria, adapted to each population.

3.1.5 *Mechanisms underlying sparsening*

Subthreshold KC activity during odor stimulation. Sharp electrode recordings (see section 3.3.6, p. 90) from KCs ($n = 29$) revealed odor-evoked, subthreshold activity made up of periodic synaptic potentials (figure 3.5A). These were locked to the LFP (figure 3.5B) and superimposed on a noisy and irregular synaptic background, away from firing threshold. Appropriate odor-KC combinations revealed reliable and time-specific EPSPs and/or action potentials. The response of the KC in figure 3.5C for example, contained a train of prominent EPSPs, late within the stimulus. One of these EPSPs led to an action potential in half of all trials with that odor. A different KC responded to the same odor with at least two reliable action potentials, at cycles 1 and 3 of the response, whether the neuron was at rest (figure 3.5D) or held depolarized by current injection. In all tested KCs, the existence, timing and reliability of these firing events were odor specific. We noted that a large component of the odor-evoked activity in KCs was inhibitory: if the KC was held depolarized by current injection, periodic hyperpolarizing potentials could be seen during a response; if the KC was held above firing threshold, odor-evoked inhibition interrupted this tonic firing (figures 3.5A and C). Odor stimulation thus also causes synaptic inhibition of KCs. Finally, the amplitude of odor-evoked EPSPs paradoxically increased when the KC was held depolarized (figure 3.5A), suggesting active membrane properties. We examined the possibility that synaptic inhibition and KC active conductances work together to make KCs coincidence detectors of PN input.

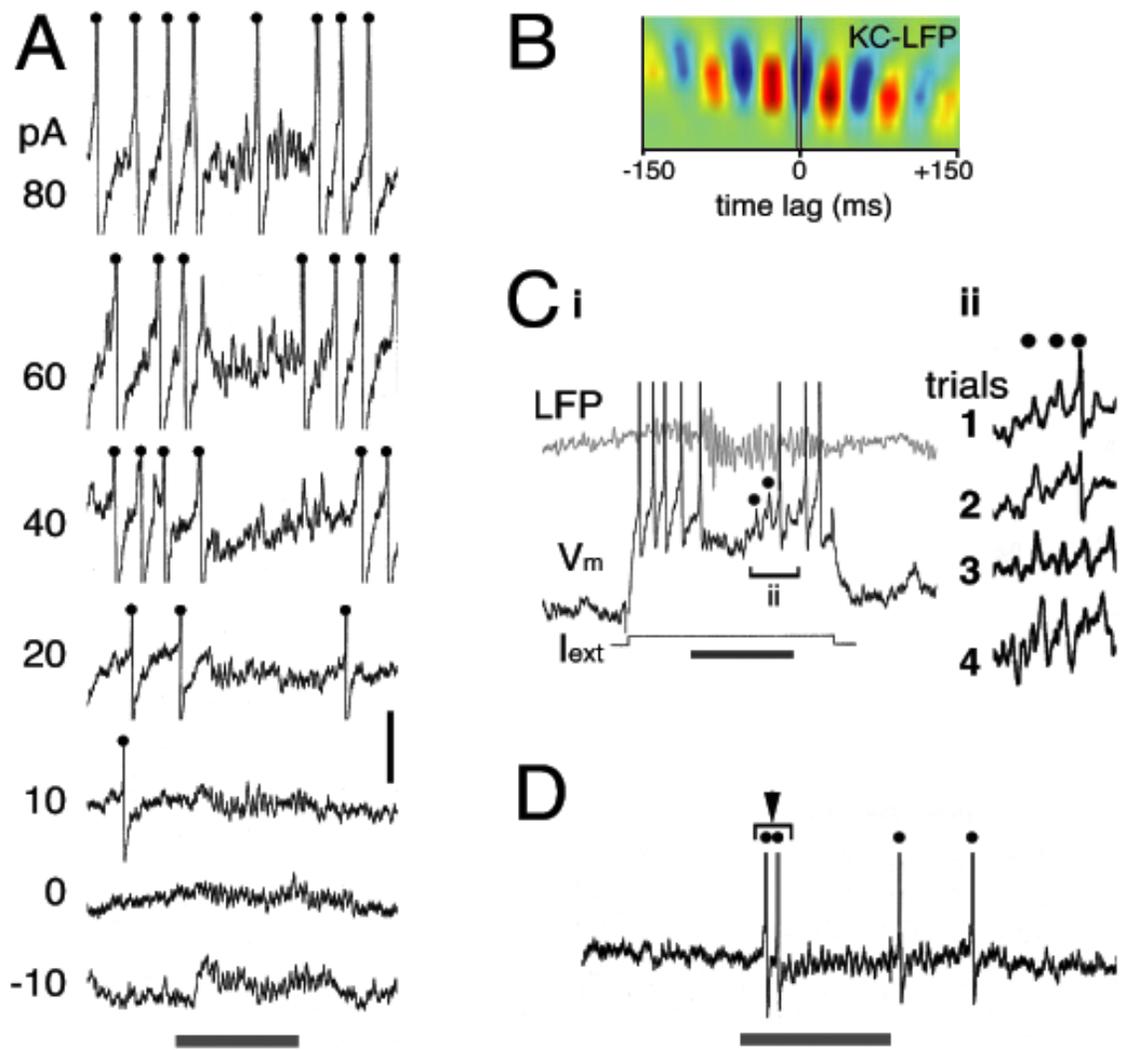


Figure 3.5. In vivo sharp-electrode intracellular records from different KCs. All action potentials are clipped. *A*, responses obtained while resting voltage set by holding currents between -10 and $+80$ pA. Horizontal bar: odor (cherry) delivery, 800 ms. This KC never produced any action potential in response to this odor at resting potential. Note oscillating membrane potential at rest (0 pA), interruption of DC-evoked firing by odor delivery ($+20$ to 80 pA traces), amplification of many discrete depolarizing potentials at most depolarized holding potentials. *B*, sliding cross-correlation of KC V_m and simultaneous LFP (different KC from *A*). Red: maxima; blue: minima. Y-axis: time, 0.5 s. Note locking of signals during odor puff. *C*, third KC recording, showing interruption of current-evoked firing by odor response and prominent, late EPSPs (●); 800 ms odor delivery (*i*). Repeated trials (1–4) show precise re-occurrence of these EPSPs during same epoch of the response; time calibration: 500 ms (*ii*). *D*, fourth KC and its spiking response to cherry odor at cycles 1 and 3. 800 ms odor pulse. Vertical calibrations: KC: 10 mV (*A*, *D*), 8 mV (*Ci*), 6 mV (*Cii*); LFP: 300 μ V, 1–40 Hz bandpass (*C*). [KC data (*A*, *C*, *D*): G. Laurent. Cross-correlogram (*B*): G. Turner.]

SOURCE OF MASKED ODOR-EVOKED INHIBITION

Because direct effects of PNs are excitatory and because locust PNs do not contain GABA (Leitch and Laurent, 1996), the source of odor-evoked KC inhibition should be downstream of PNs. In addition to sending collaterals into the mushroom body, PN axons terminate in the lateral horn (LH) (Hansson and Anton, 2000). We identified among their targets there a cluster of ~60 GABA-immunoreactive neurons (Lateral Horn Interneurons, LHIs), with direct axonal projections to the MB (see section 3.3.7, p. 91) (figure 3.6A). Intracellular staining of individual LHIs showed profuse axonal collaterals, overlapping with KC dendrites (figure 3.6B). KC dendrites receive GABAergic input (Leitch and Laurent, 1996). LHIs thus appeared well suited to be a source of the odor-evoked inhibitory inputs.

LHI RESPONSES TO ODORS

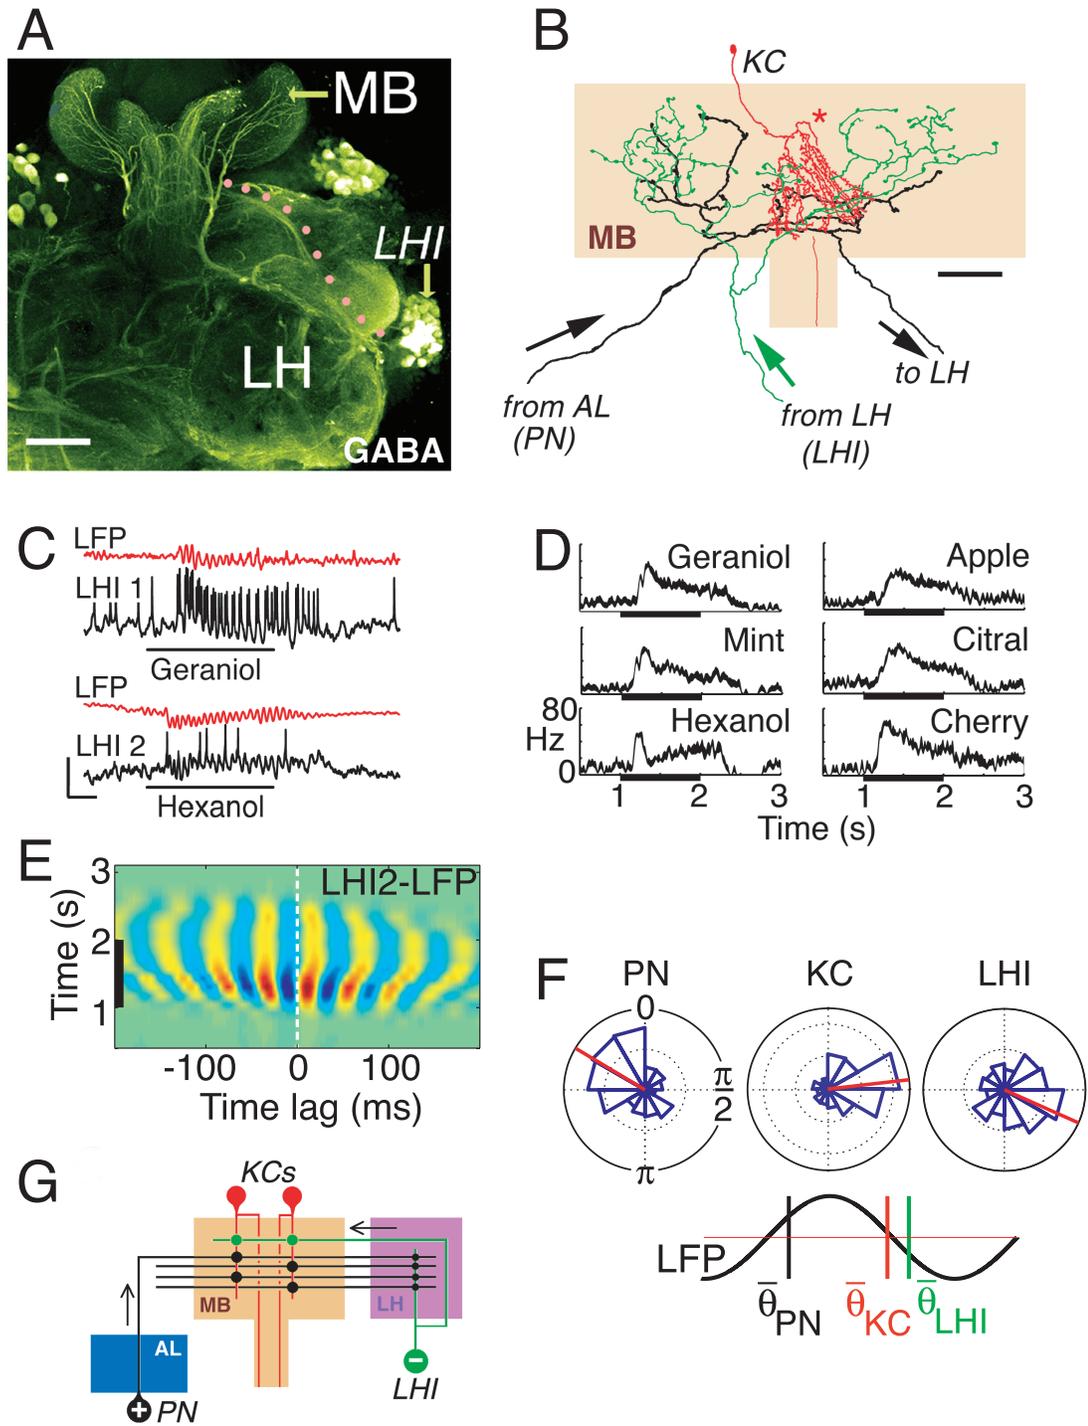
LHIs responded vigorously and reliably to odors (Fig. 4C,D). LHI membrane potential oscillated in phase with the LFP (figure 3.6E), and when sufficiently excited, LHIs fired one or a short burst of action potentials at each oscillation cycle (figure 3.6C). In each cycle, LHI mean firing-time lagged 173° behind that of PNs (figure 3.6F). LHI firing phase was independent of odor identity. Synaptic drive to KCs thus likely consists of EPSPs from PNs alternating with IPSPs from LHIs, occurring preferentially in opposite halves of each oscillation cycle. PN and LHI inputs to KCs differ in one important respect: because each KC on average receives inputs from a very small fraction of the PNs and because the firing probability and phase-locking of each PN typically evolves during a response, the probability that

many of the PNs presynaptic to a given KC fire together within the same half of one oscillation cycle is low. By contrast, individual LHIs showed sustained responses to all odors presented (figure 3.6D), consistent with the fact that 830 PNs converge onto only ~60 LHIs. Because LHI axons diverge profusely in the mushroom body (figure 3.6B), individual KCs should receive periodic input composed of consistent IPSPs, alternating with EPSPs whose total strength depends strongly on the stimulus.

SHARPENING OF KC RESPONSE TO DIRECT PN STIMULATION

We next tested more directly whether both synaptic inhibition and intrinsic active conductances assist coincidence detection in KCs. To study single EPSP-IPSP cy-

Figure 3.6 (on the next page). Feed-forward inhibition of KCs by LHIs. *A*, anti-GABA immuno-labeling (see section 3.3.7, p. 91). Cluster of ~60 reactive somata (LHI) and tract of LHI axons running to the MB (stipples). The terminals of one of these axons in the MB are shown in *B*. Calibration: 100 μm . *B*, PN axon (black) projects to the mushroom body calyx (orange) (Laurent and Naraghi, 1994) and to the lateral horn (LH). LHI (green) project to the calyx (this study). PN and LHI axons terminate on KC dendrites (red). Neurons stained by iontophoresis of cobalt hexamine (KC, PN) or neurobiotin (LHI), in separate preparations and drawn using a camera lucida. Note varicosities in LHI and PN axon collaterals. (*): KC axon. Calibration: 50 μm . *C*, representative odor-evoked responses of two LHIs and simultaneously recorded LFPs (5–40 Hz bandpass). Note membrane potential oscillations, locked to the LFP. Identity and delivery (1 s long) of stimulus indicated by black bar. Calibration: (LHI) 20 mV; (LFP) 400 μV ; 200 ms. *D*, instantaneous firing rate of LHI1 (in *C*) in response to various odors. Lower edge of profile: mean instantaneous rate, averaged across trials; profile thickness: SD. All LHIs responded to all odors tested, with response profiles that varied little across different odors. *E*, sliding cross-correlation between LFP and LHI2 traces (spikes clipped). High correlation values in hot colors, low in cold. Strong locking is present throughout the response (odor delivery: vertical bar). Lower edge of correlation stripes just precedes stimulus onset due to width of correlation window (200 ms). *F*, phase relationships between PN, KC and LHI action potentials, and LFP. Polar plots: LFP cycle maxima defined as 0 rad, minima as π rad (PNs: 3 cell-odor pairs, 388 spikes; LHIs: 17 cell-odor pairs, 2632 spikes; KCs: 18 cells, 862 spikes). Mean phases shown in red. Gridlines are scaled in intervals of 0.10 (probability per bin). Below: diagram showing LFP and mean firing phases, Q. *G*, circuit diagram. [LHI data (A–F): G. Turner. PN data (F): M. Westman. KC data (F): J. Perez-Orive.]

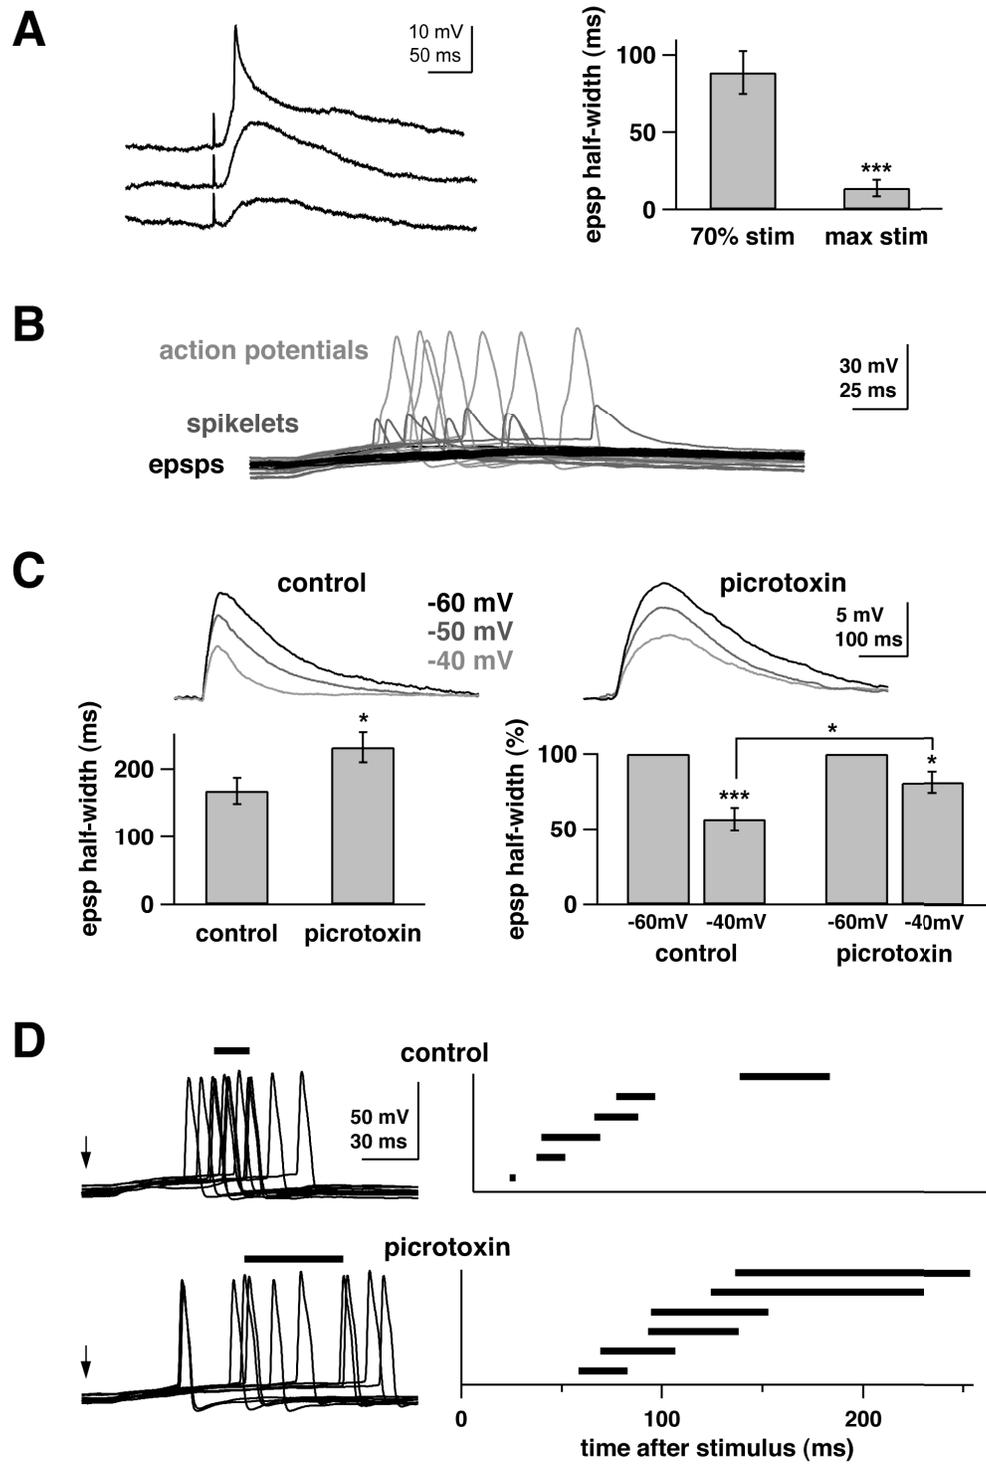


cles in isolation, we used direct electrical stimulation of PNs rather than odors. Evoked postsynaptic potentials in KCs dramatically changed shape and duration upon varying stimulus strength (figure 3.7A). At high stimulus intensities, a sharp *spikelet* rode atop the depolarizing potential, suggesting active conductances (figure 3.7A, top trace). This spikelet was not an artifact of unusually strong stimuli: when a weak stimulus was used to elicit a smaller EPSP and holding current was adjusted so that the KC was near firing threshold, spikelets could also be observed⁴ (figure 3.7B) (see section 3.3.8, p. 92; Schafer et al., 1994). Next, we tested the idea that GABAergic feed-forward inhibition also shapes PN-evoked PSPs. At voltages below spikelet threshold, EPSP shape was still strongly voltage dependent (figure 3.7C). Local injection of picrotoxin (PCT), a GABA_A-like chloride channel blocker, into the mushroom body calyx (see section 3.3.9, p. 92) broadened the EPSP and decreased the voltage-dependence of EPSP shape (figure 3.7C). This indicates that the LHI-mediated IPSP normally contributes partly, but maybe not entirely, to the shape and duration of PN-evoked EPSPs. Blocking inhibition in the calyx increased the scatter of KC-spike times following PN stimulation (figure 3.7D). LHI-mediated IPSPs thus contribute to shortening the epoch during which a KC remains depolarized after each volley of PN excitation; it could also explain why KC action potentials are so precisely phase-locked during responses to odors (figure 3.6F). Hence, the tendency of each KC to convert its excitatory input from PNs

⁴Although we have not characterized this spikelet pharmacologically, its shape and all-or-none waveform suggest the involvement of voltage-dependent conductances (possibly Na⁺ or Ca²⁺ for depolarization and K⁺ for repolarization), consistent with previous patch-clamp studies in vitro (Schafer et al., 1994)

into an action potential can be facilitated in the early phase of the compound EPSP by voltage-dependent depolarizing nonlinearities and antagonized shortly thereafter by feed-forward inhibition. The remaining voltage dependence of the EPSP after PCT injection (figure 3.7C) suggests the existence of an active repolarizing conductance. Active and synaptic properties thus both likely contribute to making KCs prefer coincident input, on a cycle-by-cycle basis.

Figure 3.7 (on the next page). KC responses to electrical stimulation of PNs. *A*, PNs were stimulated directly using an electrode placed in the AL and evoked EPSPs were recorded intracellularly from KCs. Three traces show EPSPs recorded at progressively stronger stimulus intensities (bottom to top). Note positive inflexion during rising phase of the top EPSP and sharp repolarization. Bar graph compares EPSP half-width at the maximum stimulus intensity that was still below action potential threshold vs. half-width at 70% ($\pm 5\%$) of this maximum intensity. EPSP half-width was significantly different at these two stimulus intensities ($p < 0.001$, paired *t*-test, $n = 11$ KCs). *B*, intrinsic active conductance amplifies and sharpens EPSPs near threshold. KC held near threshold with constant holding current; constant PN stimulus amplitude: successive trials elicited full-blown sodium spikes (light grey), subthreshold EPSPs (black), or intermediate *spikelet*. Sample traces collected in picrotoxin; similar spikelets were observed in control conditions (e.g., *A*). *C*, synaptic inhibition shortens KC EPSP. At progressively depolarized holding potentials, EPSP half-width significantly decreased (half-width at -40 mV significantly smaller than half-width at -60 mV, $p < 0.0005$, paired *t*-test, $n = 10$ KCs); all analyzed data below threshold for spikelet activation. After picrotoxin injection in MB, EPSPs became broader (-60 mV half-width significantly increased in picrotoxin, $p < 0.05$, *t*-test, $n = 9$). EPSP shape was less dependent on postsynaptic voltage (-40 mV half-width as percentage of -60 mV half-width significantly increased in picrotoxin, $p < 0.05$, *t*-test, $n = 9$), but was still voltage dependent ($p < 0.05$, *t*-test, $n = 9$). Sample traces from two KCs in same brain. *D*, synaptic inhibition narrows the window in which KCs can fire following PN stimulation. Stimulus intensity was adjusted to elicit an EPSP of 5–10 mV (when KC held at -60 mV), and then holding current was adjusted so that this EPSP elicited a spike on 30–60% of trials. Representative traces (left) show those sweeps that elicited spikes (arrows mark stimulus, bars mark interquartile range of spike times encompassing the difference between the 25th and the 75th percentile). Sample traces from two KCs in same brain. Group data (right) shows the interquartile range for each cell. Picrotoxin significantly increased the magnitude of the interquartile range ($p < 0.05$, *t*-test, $n = 6$ control KCs, 6 KCs in picrotoxin). *B–D*: whole-cell recordings (see section 3.3.8, p. 92). [Sharp microelectrode recordings (*A*): J. Perez-Orive. Whole-cell patch recordings (*B–D*): R. Wilson.]



3.1.6 Influence of feed-forward inhibition on KC responses to odors

If feed-forward inhibition competes with and resets the periodic excitation of KCs by PNs, antagonizing LHI-mediated inhibition should decrease KC specificity to odors. KCs recorded *in vivo* with tetrodes were tested with up to 17 odors (ten trials per odor) and re-tested immediately after PCT injection into the mushroom body (figure 3.8). PCT caused no significant change in the KC baseline-firing rate (medians: 0.018 spikes/sec after PCT vs. 0.005 before, $n = 12$ KCs, $p = 0.19$, non-parametric sign test). PCT caused a broadening of KC tuning, characterized by greatly reduced odor selectivity (figure 3.8A–C). Even in KCs that responded to none of the odors presented in controls, responses to these same odors appeared after PCT (figure 3.8A–C). Individual KCs did not become responsive to all odors, but rather, to a larger subset of all tested odors. The mean population and lifetime sparseness calculated over this KC subset was significantly decreased after PCT ($S_p = 0.70$ to 0.41 , $n = 11$ odors, $p < 0.001$, paired t -test; $S_L = 0.47$ – 0.30 , $n = 12$ KCs, $p < 0.05$, non-parametric Wilcoxon signed-rank test). Individual KC response intensity after PCT treatment was not significantly different from control (1.96 ± 0.81 spikes; PCT: 1.82 ± 0.47 spikes), but KC action potentials after PCT lost their locking to the LFP (figure 3.8D, cf. controls, figure 3.6F). This confirms earlier experiments (figure 3.7D) suggesting that LHI-mediated IPSPs normally constrain KC integration and spike timing.⁵

⁵PCT application to the mushroom body did not affect the LFP oscillations recorded there, for the principal source of these oscillations—synchronized, periodic synaptic input drive from PNs—was excitatory and cholinergic (nicotinic).

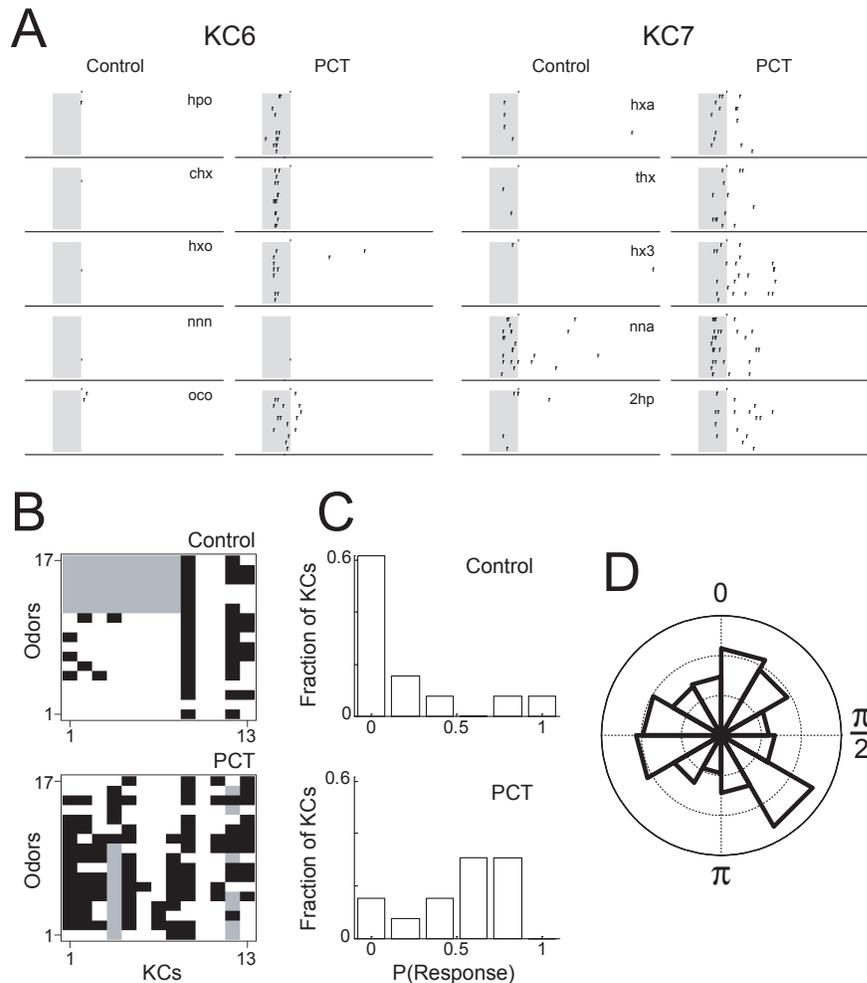


Figure 3.8. Influence of feed-forward inhibition on KC odor tuning and phase-locking (in vivo, wire tetrode recordings). *A*, two KCs (5–6) and their responses to five odors before and after local PCT injection into the MB (see section 3.3.9, p. 92). Odor pulses (shaded area): 1 s; 10 trials per odor, top to bottom (abbreviations in section 3.3.1, p. 82). *B*, comparison of KC response profiles before and after PCT. Filled: response; open: no response; grey: not tested (more odors were generally tested after PCT treatment). PCT broadened KC response tuning profiles, by causing the appearance of responses to new odors. Odors, 1 to 17: oca, hxa, thx, hx3, oco, unn, nna, 2hp, che, chx, hxo, don, nnn, 3hp, hpo, pnn, min. *C*, frequency distribution of response probabilities (across all odors tested) before and after PCT treatment ($n = 13$ KCs). Note dramatic reduction of proportion of specific KCs (leftmost bin) after PCT. Median response probabilities: 0.09 (control), 0.59 (PCT). *D*, phases of KC spikes relative to LFP, during odor-evoked responses (0 rad: max; π : min of LFP voltage). Vector strengths: 0.03 (PCT) vs. 0.41 (control, Fig 4F). Gridlines in intervals of 0.05 (probability per bin). [J. Perez-Orive]

3.2 Discussion

In the AL, individual odors are represented by a large fraction of the 830 PNs: Baseline activity is high, sparseness is low, individual PN responses are sustained. In the MB, the same odors activate a small proportion of neurons in a larger population (50,000 KCs): baseline activity is close to zero, sparseness is high, individual KC responses are rare and typically contain two action potentials only. KC action potentials thus each carry much more information than those of PNs.

HOW DOES SPARSENING ARISE?

We propose that KCs act as selective coincidence detectors on periodic PN input: Because individual KCs receive inputs from only a small fraction of all PNs, because the patterned responses of individual PNs are staggered in time and because EPSP summation by KCs occurs best within a fraction of each oscillation cycle, the conditions appropriate for bringing a KC to threshold are rarely met. During odor stimulation, each oscillation cycle contains both locked and unlocked PN spikes (Laurent et al., 1996). Periodic IPSPs, caused in KCs by LHIs whose mean firing is in antiphase with the discharge of the synchronized PNs, antagonize the action of inappropriately timed PN action potentials. When LHI-mediated inhibition is blocked, this normally antagonized excitatory drive to KCs can now summate over a longer time window: KCs lose much of their specificity. Time-locked feed-forward inhibition thus helps define very short but renewed (once per oscillation cycle) integration windows for each KCs, akin to a periodic reset, with critical consequences for KC specificity.

None of the features uncovered so far (oscillatory patterning, feed-forward inhibition, fan-in and fan-out, active properties) are unusual ones (Fricker and Miles, 2000; Galarreta and Hestrin, 2001; Pouille and Scanziani, 2001; Contreras et al., 1997; Haberly, 1990). In particular, distributed and partly overlapping projection patterns of mitral cells have been seen in rodent prepiriform cortex (Zou et al., 2001) and local feed-forward inhibitory circuits are common (Fricker and Miles, 2000; Pouille and Scanziani, 2001; Contreras et al., 1997; Haberly, 1990). Nonlinear intrinsic properties have been seen in some cortical and hippocampal cells among others and hypothesized to underlie coincidence detection (Fricker and Miles, 2000; Galarreta and Hestrin, 2001). We showed here that all these properties exist together in the same circuit and that their concerted use in the context of oscillatory activity results in a major transformation of sensory codes.

HOW COULD SPARSENING BE USEFUL?

Because the mushroom body is a likely site for the formation and retrieval of olfactory memories (Heisenberg et al., 1985; Dubnau et al., 2001; McGuire et al., 2001; Zars et al., 2000), we must ask why sparse codes might be advantageous there. While it is clear that extremely sparse codes (“grandmother” schemes, Barlow, 1969) may be undesirable because they confer sensitivity to damage and low capacity, representations carried by small subsets of neurons offer many theoretical advantages. First, overlaps between individual representations are less likely than if each representation used a large proportion of the available neurons, limiting interference between memories. This system’s memory capacity can still be very high,

because the total population size is large and sparseness is not extreme. Second, comparisons between stimulus-evoked patterns and stored memories are simpler if they invoke fewer elements. Third, representations become more synthetic or high-level. Every KC action potential compresses the signals carried by several PNs that are each potentially more informative about stimulus composition. Sparsened representations thus contain less explicit detail. This conclusion agrees with behavioral and psychophysical observations in humans, rats and insects that odor perception has a prevalent synthetic quality (Livermore and Laing, 1996; Linster and Smith, 1999; Cain and Potts, 1996).

SIGNIFICANCE FOR NEURAL CODING

Our results have implications for the understanding of neural codes. First, single-neuron responses can be exquisitely specific, extremely short (1–2 spikes only), and temporally precise (both within and across oscillation cycles). Whether response characteristics similar to those shown for KCs exist in other systems is not known; because they are so brief and specific, such response patterns are de facto hard to uncover. Studies of primate frontal and motor cortices that show very brief firing events, however, are consistent with some of our results (Abeles et al., 1993; Riehle et al., 1997). Second, subtle yet highly relevant activity patterns may go undetected with many large-scale brain-activity monitoring techniques: sparse and brief activity is unlikely to be reflected in most macroscopic signals. Yet, as we show here, this may sometimes be all there is. Finally, to measure the relevant information content of an action potential, one must know how downstream targets interpret it. For

example, we showed previously that projection neuron action potentials typically phase lock to the LFP only during certain (stimulus- and PN- specific) epochs of a response (Laurent et al., 1996). Our results indicate that KCs will be more sensitive to phase-locked PN action potentials than to those occurring closer to each LHI-mediated IPSP, whose timing is itself determined by the locked-PN population. PN spikes, therefore, are not all equally meaningful to a KC. Even in cases where firing rates are high, many spikes may be of minimal significance to a target, because improperly-timed. Here, relevance is determined by inter-neuronal correlation. Hence, deciphering brain codes requires an evaluation of these correlations and their consequences on the channeling of information. Conversely, macroscopic oscillations may indicate the existence of neural filters, whose properties will determine the interpretation one should make of a spike train.

3.3 Methods

3.3.1 Preparation and stimuli

Results were obtained from locusts (*Schistocerca americana*) in a crowded, established colony. Young adults of either sex were immobilized, with one or two antennae intact for olfactory stimulation. The brain was exposed, desheathed and superfused with locust saline, as previously described (Laurent and Naraghi, 1994). Odors were delivered by injection of a controlled volume of odorized air within a constant stream of dessicated air. Teflon tubing was used at and downstream from the mixing point to prevent odor lingering and cross-contamination. Odors

were used at 10% vapor pressure (all PNs, 85% of KCs) or 100% vapor pressure (15% of KCs), further diluted in the dessicated air stream. We used: 1-hexen-3-ol (hx3), trans-2-hexen-1-ol (thx), cis-3-hexen-1-ol (chx), 1-hexanol (hxo), 1-heptanol (hpo), 1-octanol (oco), hexanal (hxa), heptanal (hpa), octanal (oca), nonanal (nna), 3,7-dimethyl-2,6-octadiene-nitrile (don), 3-pentanone (pnn), 2-heptanone (2hp), 3-heptanone (3hp), 5-nonanone (nnn), 6-undecanone (unn), cherry (che), mint (min), geraniol (ger), vanilla (van), citral (cit), apple (app), strawberry (str), amyl acetate (ama), benzaldehyde (bnh), methyl salicylate (mts), eugenol (eug), L-carvone (lca), D-carvone (dca), dihydro-myrcenol (dhm).

3.3.2 *Tetrodes*

Two types of tetrodes were used for extracellular recordings. Silicon probes were generously provided by the University of Michigan Center for Neural Communication Technology (<http://www.engin.umich.edu/facility/cnct/>). Wire tetrodes were constructed with insulated 0.0005" and 0.0004" wire (REDIOHM wire with PAC insulation). Four strands of wire were twisted together and heated to partially melt the insulation. The tip was cut with fine scissors and each channel tip was electroplated with gold solution to reduce the impedance to between 200 and 350 k Ω at 1 kHz. The same custom-built 16-channel preamplifier and amplifier were used for both types of tetrodes. Two to four tetrodes were used simultaneously. The preamp has a unitary gain, and the amplifier gain was set to 10,000x. Because of low baseline activity and low response probability in KCs (see sections 3.1.2 and 3.1.3),

fewer KCs than PNs were usually isolated in a typical recording session. Tetrodes were placed within the AL or MB soma clusters, peripheral to the neuropils at depths between 50 and 200 μm . Cell identification was unambiguous because PNs are the only spiking neurons in the locust AL, (LNs do not produce sodium action potentials, Laurent and Davidowitz, 1994), and because all the somata located above the MB calyx belong to KCs.

3.3.3 *Extracellular data analysis*

Tetrode recordings were analyzed as described in Pouzat et al. (2002). Briefly, data from each tetrode was acquired continuously from the four channels (15 kHz/channel, 12 bit/sample), filtered (custom-built amplifiers, band-pass 0.3–6 kHz) and stored. Events were detected on all channels as voltage peaks above a pre-set threshold (usually 2.5–3.5 times each channel's signal SD). For any detected event on any channel, the same 3 ms window (each containing 45 samples) centered on that peak was extracted from each one of the four channels in a tetrode. Each event was then represented as a 180-dimensional vector (4×45 samples). Noise properties for the recording were estimated from all the recording segments between detected events, by computing the auto- and cross- correlations of all four channels. A noise covariance matrix was computed and used for noise whitening. Events were then clustered using a modification of the expectation maximization algorithm. Because of noise whitening, clusters consisting of, and only of, all the spikes from a single source should form a Gaussian ($\text{SD} = 1$) distribution in 180-dimensional space. This

property enabled us to perform several statistical tests to select only units that met rigorous quantitative criteria of isolation (figure 3.9).

3.3.4 Responses

Defining what constitutes a response quantitatively and equally accurately for PNs and KCs requires careful consideration. For example, a conventional mean firing rate measure applied to the entire response period is not appropriate, because PN responses are patterned; a typical PN response, such as one composed of subsequent excitatory and inhibitory epochs, often produces a mean rate no different from baseline, and yet clearly constitutes an odor-specific response; reliability across trials thus needs to be taken into account. In addition, PNs and KCs have very different baseline firing statistics, implying that response criteria based on a change from baseline might not apply equally well to both populations. We thus analyzed the data using a variety of methods and display, in our paper, the results of one (Method A), applied identically to KCs and PNs. The analyses using other methods, summarized in table 3.1 and figure 3.10, yielded nearly identical results. Our methods are as follows. First, for all methods, we defined two response windows: short (0–1.4 s) and long (0–3 s after stimulus onset), with stimulus on for 1 s in all cases. Method A used a 3 s window. Second (Method A), a PN or KC was classified as responding during either window if its firing behavior during the window met two independent criteria of response amplitude and reliability:

1. **AMPLITUDE:** The neuron's firing rate (measured in successive 200 ms bins,

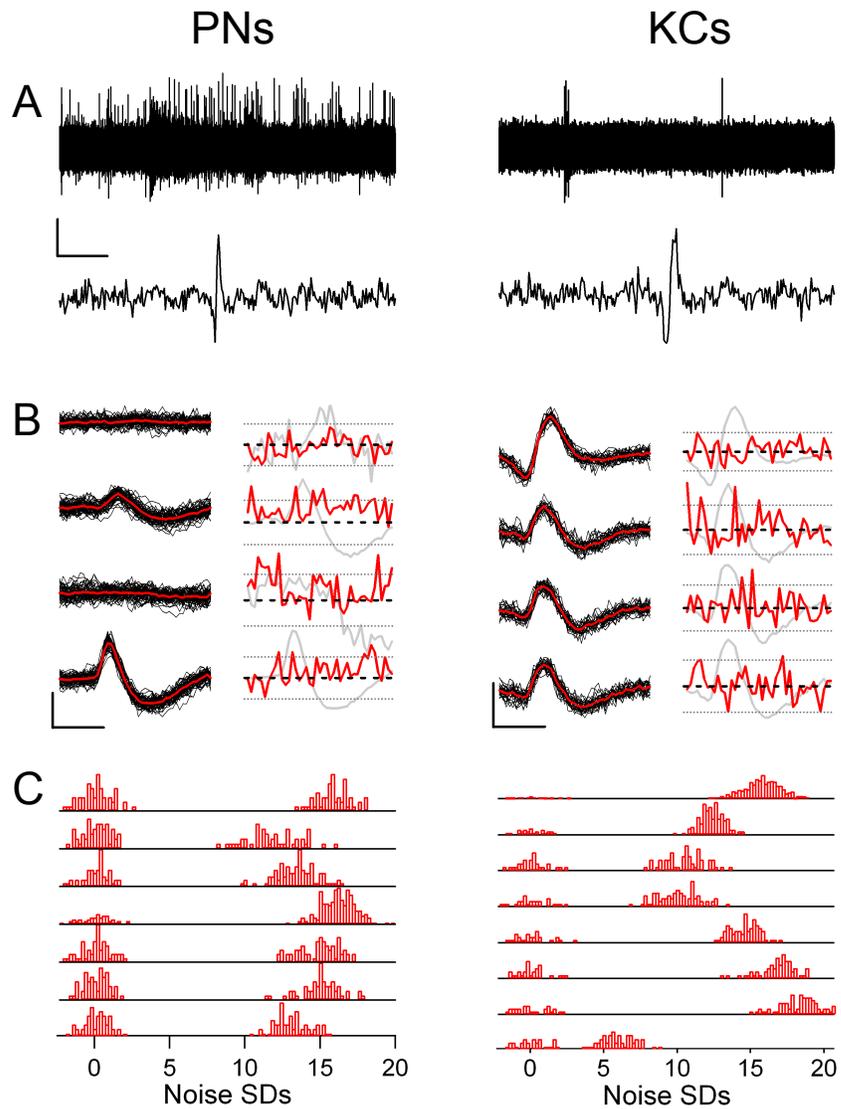


Figure 3.9. Extracellular tetraode recordings and spike-sorting. *A*, raw data traces with PN action potentials recorded in the AL (*left*), and KC action potentials recorded in the soma layer of the MB (*right*). Calibrations: 50 μV , 3 s (top traces), 3 ms (bottom traces). *B*, examples of two clusters: PN (left panel) and KC (right panel). In each panel the traces on the left show the superimposed events classified for that cluster (black) for each of the four tetraode channels, together with the average waveform (red). Calibration: 100 μV , 1 ms. Two of the statistical tests used to evaluate the isolation of the cells in the model are shown in *B* and in *C*: on the right side of each panel in *B* is the variance around the mean for each of the four channels, together with 95% confidence intervals which are based on the noise model. *C*, projection tests in which each pair of clusters in the model in 180-dimensional space is projected onto the line connecting the cluster centers so as to evaluate their degree of isolation. All cluster centers are separated by at least five times the noise SD. All analyzed data were selected on these separation criteria. [PN data: O. Mazor and S. Cassenaer. KC data: J. Perez-Orive.]

averaged across all trials) had to exceed n SDs of the mean baseline rate in at least one bin within the response window. Baseline rate was measured for each cell-odor pair over a period of 3 to 5 s preceding stimulus onset and over all trials with that odor. We explored values of n from 2 to 4. If n was low (e.g., $n = 2$ SDs) the rate of false responses detected in PNs prior to stimulation was unacceptably high (>35%). If n was high ($n = 4$ SDs), the proportion of missed responses (as judged by visual inspection of PN rasters and PSTHs) during odor presentation was unacceptably high (>10%). Values of n of 3 or 3.5 gave low rates of both false positives (during baseline) and false negatives (during stimulation) in PNs. Values of n between (and including) 2 and 4 made no significant difference with KCs. We show the results with $n = 3.5$ (Method A, figure 3.4); those obtained with other values of n are summarized in table 3.1.

2. **RELIABILITY:** To ensure that responses detected were reliable even at low firing rates, we required that more than half of all trials with each odor contain at least one spike during the response window. We also analyzed the same data sets using different criteria for PNs and KCs, each adapted to each population's baseline firing statistics. Despite this difference, the results (table 3.1, figure 3.7) are nearly identical to those shown in figure 3.4.

3.3.5 *Sparseness*

Data were analyzed using Matlab and Igor. The sparseness measures are taken from Willmore and Tolhurst (2001), Rolls and Tovee (1995), and Vinje and Gallant

Method		A	B	C	D	E	F	G	H
	<i>Threshold</i>	3.5 SD	2 SD	3 SD	4 SD	3.5 SD			
	<i>Window</i>	3 s	3 s	3 s	3 s	1.4 s	3 s	3 s	3 s
PN	P(Resp.)	0.64	0.73	0.68	0.61	0.51	0.64	0.65	-
	False Pos. (%)	2.23	35.98	6.16	0.80	0.89	2.14	3.57	-
	Overlap (%)	-	90.71	96.34	96.79	87.41	99.73	99.55	-
KC	P(Resp.)	0.11	0.12	0.12	0.11	0.09	0.11	-	0.11
	False Pos. (%)	0.09	1.46	0.18	0.00	0.00	0.09	-	0.00
	Overlap (%)	-	99.27	99.82	99.91	97.46	100.00	-	99.46

Table 3.1. Quantitative comparison of different methods of response detection. For each method, three statistics are computed for PNs and KCs. Response probability (P(Resp.)) indicates the probability of a detected response, computed over all cell-odor pairs. The false positives value (False Pos.) is the percentage of responses detected when the method was applied to a window of baseline activity prior to odor onset (computed for all cell-odor pairs). The final statistic (Overlap), is a measure of similarity between a particular method and Method A (see section 3.3.4, p. 85), defined as the percentage of cell-odor pairs for which the two methods either both detected or both did not detect a response. Methods B–D are identical to Method A, but use a different response amplitude threshold, ranging from 2 SDs to 4 SDs above baseline. Method E is the same as Method A, but uses only a 1.4 s response window (0–1.4 s after odor onset). Method F is based on Method A, but it uses a different reliability criterion that adapts to the cell’s baseline statistics. In this method, an odor response was reliable if more than half of all trials contained at least one 200 ms bin with a spike count higher than a threshold, specified as 1 SD above the mean baseline rate. Methods G and H are the methods of response detection for PNs and KCs, respectively, described in figure 3.10. [PN data: O. Mazor and S. Cassenaer. KC data: J. Perez-Orive.]

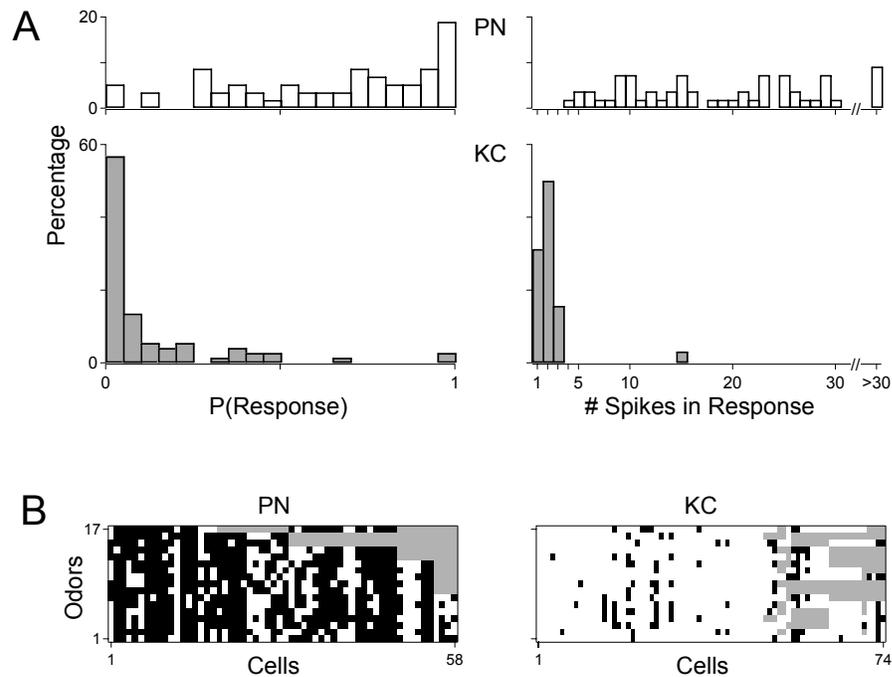


Figure 3.10. Population responses and sparseness across PNs and KCs, calculated using different criteria with PNs and KCs, for determining whether a neuron responded: A PN was qualified as responding during the 3 s following odor onset, if its firing rate increased to above 3.5 SDs of the pre-odor baseline rate (measured by a PSTH with 200 ms non-overlapping bins). In contrast, a KC response occurred when over 50% of individual trials for a particular odor showed an increase from baseline activity anywhere in the 3 s window. An increase in activity was defined as at least one 200 ms bin with a spike count higher than 3 SDs above baseline (computed from the pre-odor period over all trials). *A left*, histograms displaying PN and KC response probability distributions. Response probabilities measured across all odors tested. Note opposite skews in KC and PN distributions. *A right*, histograms displaying distributions of spike numbers in a response. Spike counts were computed only from cell-odor pairs with a significant excitatory response during the analysis window. *B*, excitatory responses (filled boxes) of individual PNs and KCs. Open squares denote inhibitory response (PNs only) or absence of a response (see figure 3.4B legend for odors). [PN data: O. Mazor and S. Cassenaer. KC data: J. Perez-Orive.]

(2000). In brief,

$$S_P = \frac{1 - \left(\frac{\sum_{j=1}^N r_j / N}{\sum_{j=1}^N r_j^2 / N} \right)}{1 - 1/N},$$

where N is the number of units and r_j is the response of unit j . Lifetime sparseness, S_L , is calculated in the same way, except that index j now corresponds to each odor and N to the total number of odors tested with each cell. Analog response intensities for a given cell-odor pair were computed by first segmenting the recording into 200 ms bins and computing the mean spike count in each bin, averaged over all trials with that odor. We then subtracted from all bin measures within the analysis window (1.4 or 3 s), the mean baseline rate. All so-calculated values greater than 0 over the window (7 or 15 bins) were then added. S_P and S_L vary between 0 and 1 (1 = most sparse).

3.3.6 *Sharp pipette recordings and staining*

Sharp electrode recordings of KCs (figures 3.5 and 3.7A) were made with borosilicate glass micropipettes ($R_{DC} > 300 \text{ M}\Omega$) filled with 0.2 or 0.5M K^+ -acetate or patch-electrode solution (see section 3.3.8). KC input resistance at the soma was usually around 1 G Ω . Intracellular recordings of LHIs (soma or dendritic impalement, figure 3.6) were made with borosilicate glass micropipettes filled with 0.5M K^+ -acetate (R_{DC} :100–300 M Ω) or wire tetrodes. Intracellular staining of LHIs was carried out by iontophoretic injection of 2% neurobiotin in 0.5M K^+ -acetate (0.5 s current pulses of -2.5 to -3.5 nA at 1 Hz for 30–60 min). Injected neurons were visualized in whole mounts using a diaminobenzidine-based chromogenic reac-

tion (Wicklein and Strausfeld, 2000). Local field potentials were always recorded in the mushroom body calyx, using saline-filled patch pipettes (R_{DC} : 2–15 M Ω) or wire tetrodes. Electrical stimulation of PNs was carried out in the AL using 25 μ m bipolar tungsten wires and a WPI stimulus isolator.

3.3.7 Immunocytochemistry

Partially desheathed locust brains were fixed for one hour in 5% formaldehyde, desheathed and washed for 20 h in PBS. Brains were then dehydrated through an ethanol series, placed in propylene oxide for 20 min, rehydrated and then agitated for five hours in PBS containing 5% triton and 0.5% bovine serum albumin (PBS 5% T 0.5% BSA). They were then washed for 30 min in PBS 0.5% T 0.5% BSA, and transferred to fresh PBS 0.5% T 0.5% BSA containing anti-GABA at 1:100 dilution, or, for negative control, to PBS 0.5% T 0.5% BSA lacking primary antibody. After incubation at 4°C for six days, brains were washed for two hours in PBS at room temperature and transferred to PBS 0.5% T 0.5% BSA containing fluorescein isothiocyanate-conjugated goat anti-rabbit IgG at 1:20 dilution and incubated at 4°C for four days. They were then washed for 30 min in PBS, dehydrated through ethanol series, cleared in methyl salicylate and examined by confocal laser scanning microscopy. Figure 3.6A is a projection along the z-axis of a stack of 30 optical slices each 2.7 μ m thick, constructed using the public domain *ImageJ* program (<http://rsb.info.nih.gov/ij/>). Negative control brains showed diffuse background staining.

3.3.8 *Patch-clamp recordings*

Whole-cell patch-clamp recordings from KCs were obtained in a semi-reduced preparation. After the brain was exposed, it was removed from the head with antenna and eyes still attached, placed on a glass coverslip in a custom-built chamber, and immobilized using insect pins placed in the eyes. The brain was then desheathed. Recordings were obtained from KC somata under visual control using a microscope with IR-DIC imaging. Patch pipettes (5–6 M Ω) were filled with a solution of (in mM): K gluconate 185, HEPES 10, EGTA 1, MgATP 4, Na3GTP 0.5 (335 mOsm, pH 7.2). Glucose (10 mM) was substituted for an equimolar amount of sucrose in the external saline solution, and the saline was bubbled continuously with O₂. Hyperpolarizing current injections (10 pA) were used to continually measure intrinsic membrane properties, and the cell was accepted for recording as long as $R_{input} > 1 \text{ G}\Omega$ and $R_{access} < 40 \text{ M}\Omega$. Data was acquired on an Axopatch 1D amplifier at 10 kHz and filtered at 5 kHz. Note: In whole-cell current-clamp mode, typical EPSP duration in controls at -60 mV (figure 3.7B) was about twice that observed with sharp electrodes (figure 3.7A).

3.3.9 *Picrotoxin injections*

Patch pipettes were back-filled with a solution containing 1.67 mM picrotoxin and 0.3% Fast Green. After the pipette was introduced into the MB calyx (dendritic region of the MB), a pneumatic pico-pump (WPI) was used to apply a series of four to nine 100 ms, 10 psi pressure pulses. Each pulse injected $\sim 1 \text{ pL}$ of solution

(as measured by previous injection into a drop of oil). Injected solution remained exclusively localized to calyx, as verified by dispersal of Fast Green.

3.4 Acknowledgments

The work presented in this chapter is the product of a strong collaboration between Ofer Mazor, Javier Perez-Orive, Stijn Cassenaer, Glenn C. Turner, Rachel I. Wilson, and Gilles Laurent. O. Mazor and S. Cassenaer performed the extracellular PN recordings (figures 3.2, 3.3, 3.4, 3.9, and 3.10). O. Mazor, J. Perez-Orive and S. Cassenaer performed the quantitative comparison between PN and KC odor responses (figures 3.4, 3.10, and table 3.1). Sparseness calculations (figure 3.4) were computed by O. Mazor. Sharp electrode intracellular KC recordings were performed by J. Perez-Orive (figures 3.2, 3.3, 3.4, 3.6, 3.7, 3.8, 3.9, 3.10) and Gilles Laurent (figure 3.5). R. Wilson conducted KC whole-cell patch clamp recordings (figure 3.7). G. Turner identified and characterized the physiological and anatomical properties of the LHIs (figure 3.6), and computed the KC-LFP cross correlogram (figure 3.5).

This chapter is reprinted from *Science*, vol. 297, J. Perez-Orive, O. Mazor, G. C. Turner, S. Cassenaer, R. I. Wilson, and G. Laurent, "Oscillations and sparsening of odor representations in the mushroom body", pp 359–65, 2002, with permission.

CHAPTER 4

Projection Neuron Population Activity: Dynamics and Coding

WHETHER NEURAL INFORMATION IS CONTAINED in the precise timing of spikes and whether neurons can decode this information, are hotly debated questions in neuroscience. Previous results (Laurent and Davidowitz, 1994; Wehr and Laurent, 1996) have demonstrated that odor information is encoded in the dynamic responses of projection neurons (PNs) in the locust antennal lobe. Recent work (Perez-Orive et al., 2002, 2004) has described a detailed mechanism for the readout of this information by the Kenyon cells (KCs), direct synaptic targets of the PNs. This study examines the detailed statistical properties of the PN population response. It quantifies the speed of the dynamics and the separation of odor representations at every time point. Finally, we highlight those components of the PN population response that are most likely to underlie the KCs response.

4.1 Results

4.1.1 *Odor evoked dynamics*

In this study, we set out to characterize and quantify the responses to odors of the projection neuron population in the locust antennal lobe. Previous studies (Laurent and Naraghi, 1994; Laurent and Davidowitz, 1994; Wehr and Laurent, 1996) revealed two general features of the odor responses of individual PNs. The first is an odor- and PN-specific pattern of excitation and inhibition, with modulations on a timescale of hundred of milliseconds (slow dynamics). The second is the tendency of PN spikes to phase lock to the odor-evoked ~ 20 Hz LFP oscillation (fast dynamics). Here we explore the progression of the slow dynamics over time by measuring population responses to odor stimuli of varying durations.

PNs were recorded extracellularly in vivo using tetrodes. Population data were assembled by combining sets of simultaneously recorded PNs across experiments, as previously described (Stopfer et al., 2003). PNs are the only neurons of the locust antennal lobe producing sodium spikes. Their identification is thus without ambiguity.

Figure 4.1A shows the response of one PN to 1 s and 10 s long pulses of the same odor over 10 trials. Figure 4.1B shows the PSTH constructed from these 10 trials for this and three other PNs. As described previously (Laurent and Davidowitz, 1994; Wehr and Laurent, 1996), these individual PNs respond to a 1 s odor pulse with a period of slow dynamics lasting ~ 2 – 3 seconds. When the same odor is presented for 10 s, however, the slow dynamics do not last for the entire duration of the odor

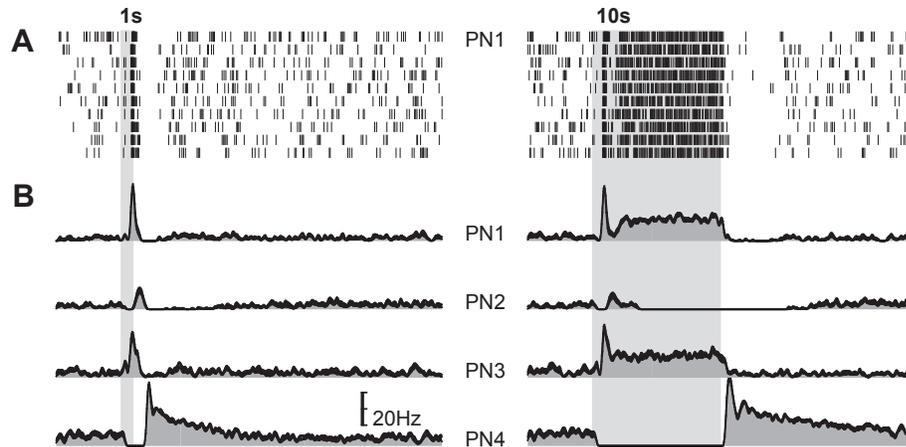


Figure 4.1. Projection neurons (PNs) respond to prolonged odor presentations with constant firing. *A*, ten single trial responses of one PN to a 1 s and 10 s presentation of the same odor. Each row corresponds to one trial, each tick represents the timing of one action potential from the cell. Trials of the two odor durations were presented interlaced. *B*, post stimulus time histograms (PSTHs) of four different PNs in response to a 1 s and a 10 s odor presentation. PSTHs are averages of ten single-trial responses, smoothed with a 100 ms Gaussian filter. Width of trace represents \pm SE. Top trace is the same PN response from part (A). Note that in all cases, the response to the 10 s stimulus reaches a constant firing rate after \sim 2–3 s.

pulse. Instead, after 1–2 s of modulated activity (initially identical to that evoked by shorter pulses), PN responses reach a noisy steady state, different across PNs and across odors. This steady state is often 0 (i.e., lower than baseline activity). At odor pulse offset, dynamics resume, such that PNs inhibited at steady state return to baseline levels (either monotonically, e.g., PN2, figure 4.1B, or after an overshoot of rebound excitation, e.g., PN4, figure 4.1B) and PNs excited at steady state return to baseline, often after some short period of inhibited firing (e.g., PN1, figures 4.1A and B).

To quantify this behavior across the PN population, we measured the statistics of PN firing throughout odor responses of 0.3 s, 1 s, and 3 s durations. We chose 3 s as the longest pulse because steady state was always reached by that time. Data were

taken from 99 PNs, stimulated with 5 odors (hexanol, cis-3-hexen-1-ol, 1-octanol, citral, methyl salicylate) diluted in mineral oil to 1% vapor pressure. Responses were computed from 9 trials per stimulus condition, representing a total of 4455 trials (99 PNs \times 5 odors \times 9 trials). The baseline PN firing rate calculated with this dataset was 2.5–3 spikes/sec. Just following odor onset and odor offset, the mean firing rate rose to \sim 4 spikes/sec. Between onset and offset, the mean PN rate was in between baseline and this peak (figure 4.2A).

The odor responses of individual PNs (e.g., figure 4.1B) indicated that specific PNs may be strongly excited or inhibited at certain times during an odor response. Mean firing rates (figure 4.2A) thus, fail to characterize the full diversity and range of activation patterns across the PN population. To address this issue, we measured the response of each PN during and around the odor trial. Response was measured as follows: each trial was divided into successive 50 ms time bins, aligned at odor onset. The number of spikes produced by each PN in each time bin was measured, and a PN was described as responding during time bin i if it produced at least one action potential in that bin over 6 or more of the 9 trials tested. This measure ensured that PN responsiveness was essentially 0 prior to odor onset (PN firing at baseline is not locked to the stimulus and thus unreliable from trial to trial) (figure 4.2B). This measure detected reliable PN firing (with 50 ms resolution) evoked by odors. PN responses measured this way are shown in figure 4.2B: each dot indicated for each PN (each line) whether this PN responded during the corresponding time bin. This is therefore not directly a measure of firing rate, but a measure of reliability.

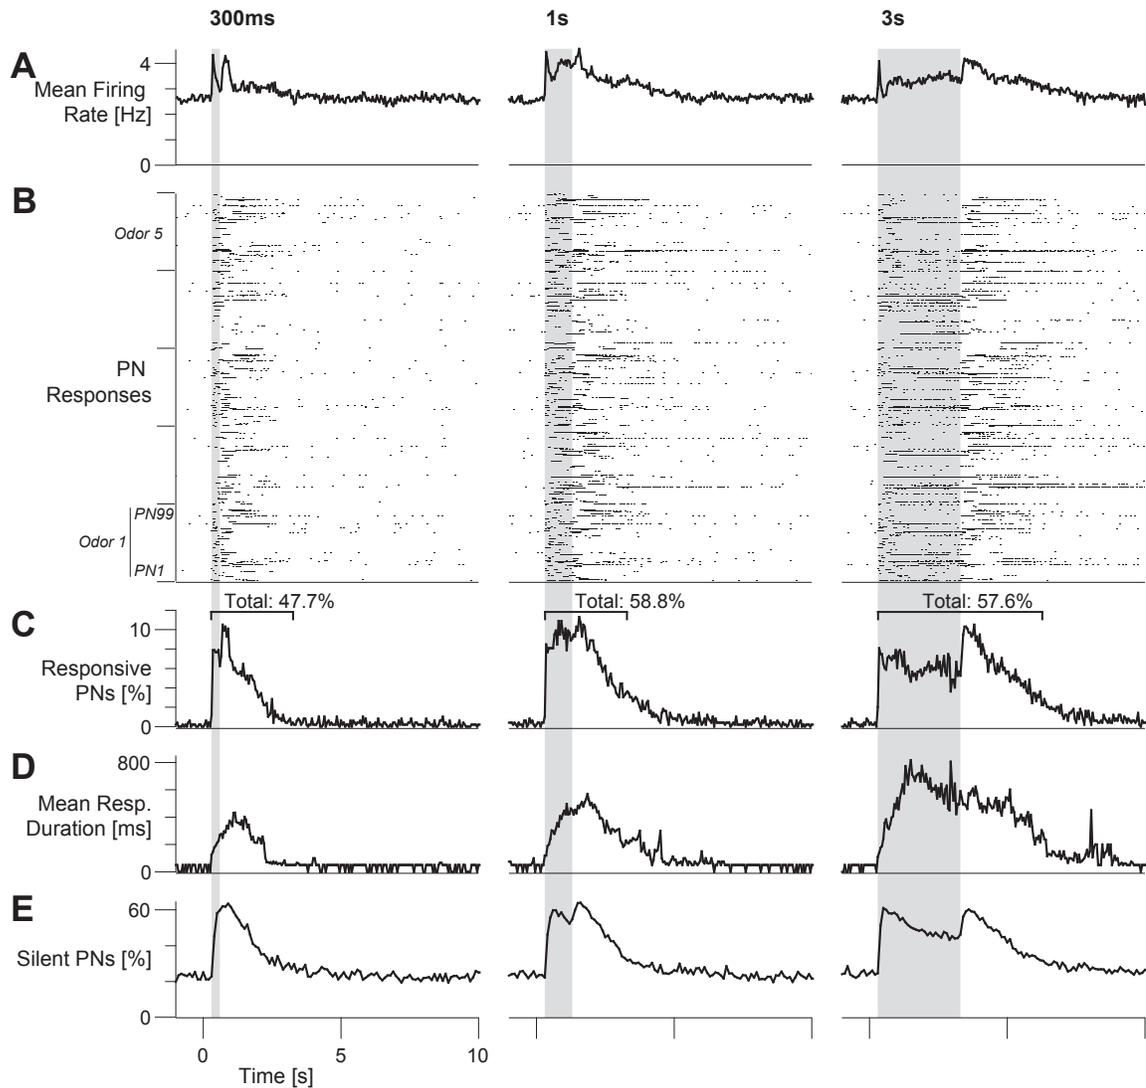


Figure 4.2. PN responses consist of small, time-varying subpopulations of responsive PNs and a large fraction of silent PNs. *A*, the mean PN firing rate over all odors for three different odor durations. *B*, raster plot for all PNs and all odors, indicating time bins when PNs responded. A response occurs if a PN produced at least one action potential in a 50 ms time bin across 6 or more of the 9 trials tested. *C*, the total percentage of responsive PN-odor pairs. *D*, the mean response duration for all PNs responding at a particular time bin. *E*, the total percentage of silent PN-odor pairs. A PN is defined as silent during a 100 ms time bin if it fired no spikes in all 9 stimulus trials. All statistics were computed from 99 PNs, each presented with 5 different odors at 3 different durations (300 ms, 1 s, 3 s). Silent PNs were computed using consecutive, non-overlapping 100 ms bins, all other statistics were computed using consecutive, non-overlapping 50 ms bins.

The percentage of responding PNs during each time bin is plotted in figure 4.2C. At rest fewer than 1% of all PNs produce action potentials at a rate high enough to be detected as responding by our measure. Following odor onset, however, the percentage of responding PNs immediately rises to between 8 and 10%. For each odor pulse condition, this level of instantaneous activity persists until about one second following odor offset, when it gradually returns to its baseline level over several seconds. This analysis revealed that, during each 50 ms bin following odor onset, no more than 10% of the PNs are reliably active, and that about 60% of the PNs are completely silent (figures 4.2C and E). Of all the spikes produced during each time bin, however, 55% are due to the responding PNs, i.e., to only 10% of all PNs. The remaining 30% of PNs were thus less reliable, or firing randomly, unaffected by the odor. This shows that during a single odor trial, more than half of all spikes produced in each time bin originate from only 10% of all PNs, but that the action of these spikes may be combined with that of other spikes produced less reliably from trial to trial.

Because individual PNs are active over only a fraction of the total odor response, the identities of responding PNs change from time bin to time bin. We thus calculated the cumulative proportion of responding PNs during an odor response (over 3 seconds for 0.3 and 1 s pulses, over 6 seconds for 3 s pulses). We found that, on average, 50–60% of PNs (figure 4.2C) can be described as responding at some time over the entire population response duration. This establishes that odor representations are broadly distributed across the PN population, but because only 10% of

PNs are co-active during any 50 ms time bin, activity must evolve across the PN population.

Accordingly, we measured the mean response duration of all responding PNs as a function of time during odor responses (figure 4.2D). For example, the mean response duration during the first second of the response to a 3 s odor pulse is ~300 ms. One second later, the mean response duration is ~600 ms, even though the percentage of responding PNs per time bin is about the same (figure 4.2C). This suggests that the speed at which PNs are being replaced slows down as the response evolves.

Just as there are periods of reliable PN activity (responses), there are periods of reliable odor-evoked inhibition. We measured the percentage of PNs that fired no action potentials within each 100 ms time bin (across 9 trials). At rest, this value was 23.2%. Just after odor onset and offset the population tripled to about 60% (figure 4.2E). In the middle of the 3 s odor presentation, the percentage fell to ~50%.

Taken together, these data show that, unlike during baseline, each time bin of the odor response is characterized by a large fraction (~60%) of completely silenced PNs and a small and evolving subset of reliably responding PNs. The remaining projection neurons are only moderately active, spiking occasionally, but not with enough consistency to meet our response criterion. Nevertheless, when the entire duration of the response is considered, about half of the PNs respond during at least one 50 ms time bin and over a few hundred ms on average and therefore

contribute to the population representation.

4.1.2 *Spatio-temporal patterns as trajectories in PN space*

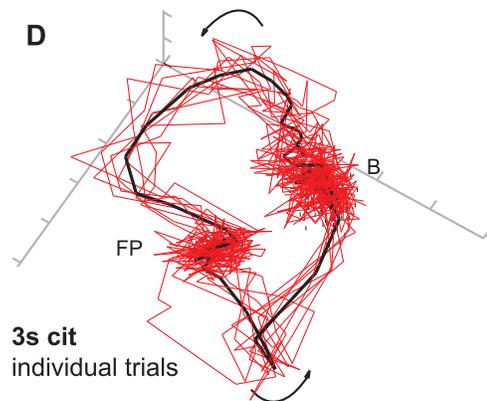
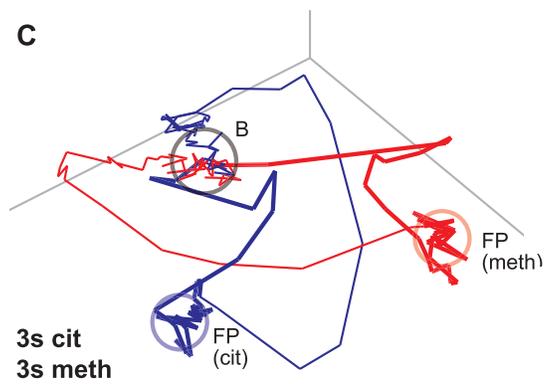
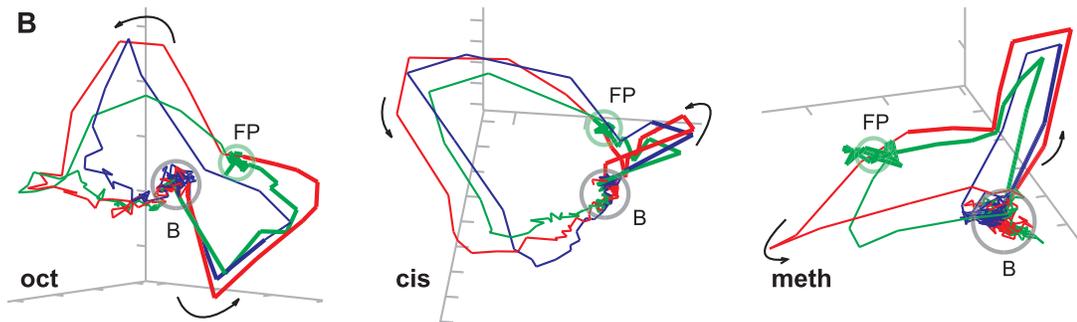
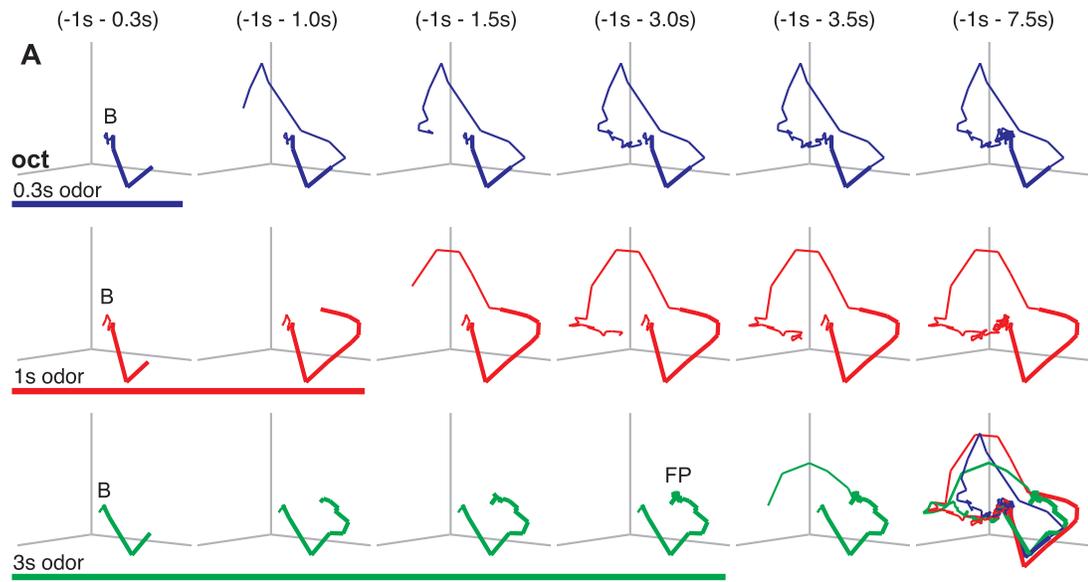
The locust antennal lobe contains approximately 800 PNs. The activity of this population of PNs at a given time can be represented as a point in an 800-dimensional space, where each dimension represents the firing rate of one of the 800 PNs at that time. When an odor is presented, subsets of PNs are activated in a PN- and odor-specific manner, as described above. If the instantaneous firing rates of the PNs are measured over short time bins (e.g., 50 ms), the state of the PN populations should describe odor-specific trajectories, beginning and ending at a point (or collection of points) that represents baseline activity. The following analysis examines the structure and correlation of these trajectories for different odors and stimulus durations, using data gathered from 99 PNs recorded from 10 animals. To allow visualization, the 99-dimensional data were analyzed with PCA and projected on to the first three principal components (figure 4.3).

Figure 4.3A shows the trajectories representing the evolution of the PN population response to three different stimulus durations (300 ms, 1 s, and 3 s) of the same odor (1-octanol). Each point along the trajectory represents the instantaneous PN activity vector averaged over 9 trials (trials 2–10) (calculated here over 100 ms bins, over the entire dataset). The thick segments of the trajectories represent the epochs during which the stimulus was present. The thin segments represent the evolution of the population response after odor termination, and the relaxation

to baseline. These plots illustrate several important points that will be examined quantitatively later: First, the initial excursions following odor onsets overlapped each other (with some noise, see below), consistent with the fact that the PN population did not know in advance the duration of the ongoing stimulus. Second, the relaxation to baseline after odor offset followed trajectories different from those defined at odor onset: the system does not retrace its steps back to rest. Third, the system reached a fixed point only for odor durations longer than about 2 s (4th frame of 3 s odor series, figure 4.3A).

The trajectories defined by the system over 7.5 seconds following each of the three pulse durations are superimposed in figure 4.3B (oct). While, as noted above, the original excursions overlapped well, the return paths to baseline were also quite similar, despite differences in past histories. This can be seen quite clearly for

Figure 4.3 (on the next page). Visualization of PN population odor responses. The activity of the population of PNs at a given time is represented as a point in 99-dimensional space, where each dimension represents the firing rate of one of the 99 PNs at that time. To allow visualization, the 99-dimensional data were analyzed with PCA and projected on to the first three principal components. *A*, the population response to odorant 1-octanol at three different odor durations (blue: 300 ms, red: 1 s, green: 3 s). Activity at successive time points are joined by a line to display the full odor-evoked trajectory. The thick segments of the trajectories represent the epochs during which the stimulus was present. The thin segments represent the evolution of the population response after odor termination, and the relaxation to baseline. All subplots use the same three principal components (PCs). *B*, comparison between response trajectories to three different stimulus durations for three different odors (1-octanol, citral, and methyl salicylate). Principal components were computed separately for each odor. *C*, odor trajectories (3 s duration) for two different odors (citral and methyl salicylate) plotted in the same PC axes (one set of PCs was computed for all point of all three odor trajectories). The fixed point for each odor is outlined by a circle. Note that the trajectories for all three odors (including fixed points) are separate and non-overlapping. *D*, example of trial-to-trial variability. The trajectories for all nine trials of the 3 s citral presentation (red traces) are plotted, along with their mean (black trace), in their first three PC axes.



the trajectories defined for two other odors (cis, meth, figure 4.3B).¹ This suggests that off-transients (i.e., segments of trajectories defined after odor offset) contain odor-specific information independent of stimulus duration.

Fifth, the fixed points reached by the PN population differed across odors (figure 4.3C) and thus contained stimulus related information. We observed, however, that the fixed points (as estimated from projections in 3-PC space) were not necessarily the points of greatest inter-trajectory separation; greater separations usually occurred during the transient (on and off) phases, when dynamics are seen in PN response patterns. Finally, we calculated single-trial PN vectors and superimposed the corresponding trajectories computed for all nine trials with one odor (citril, figure 4.3D). While illustrating the reliability of the odor-evoked trajectories, this plot also shows that the inter-trial distances between odor-evoked trajectories were no greater during the transient phases than they were at steady state (fixed points corresponding to odor or to baseline).

To quantify these qualitative observations, we calculated the mean state of the system at rest (by averaging baseline PN activity vectors, calculated in the full 99-dimensional space, over all trials and all 50 ms long time bins prior to the odor presentation); we then measured the Euclidian distance from that mean baseline PN vector to each point along each odor trajectory. The mean value of this distance (averaged over all five odors presented) is plotted in figure 4.4A. In figure 4.4B, we plot the mean inter-trajectory distance at each time point around

¹Note that the cross-over points in the trajectories are projection artifacts: the system rarely visited the same point for two different odors.

and during the odor response, averaged across all ten pairwise combinations of odor-specific trajectories. Taken together, these two figures show that the odor-evoked PN population activity is most different from baseline and from other odor responses just after the onset or the offset of the odor pulse. Both inter-trajectory distance and distance to baseline peaked around 100 ms after odor onset. At odor offset, both distance measures peaked after about 250–300 ms (250 ms for 300 ms and 3 s odor durations, 300 ms for 1 s odor pulse). In the time between the onset and offset peaks, both distance measures decreased in magnitude, but remained above baseline, indicating that odors could still be identified. In the trials long enough to allow the odor-specific fixed points to be reached (3 s trials), inter-trajectory distances were, as suggested by data in figure 4.3, greater during either transient phases than between the fixed points (figure 4.4B).

The trajectories corresponding to individual trials showed some variability around the mean trajectory (figure 4.3D). To test whether the inter-odor distances were significant relative to this inter-trial noise, we measured mean trial-to-trial distances, both within and across odor trial groups (figure 4.4C). We observed that, for all odor pulse durations inter-odor distances rose substantially above baseline levels, with peaks once again during the transient phases of the population response. Within-odor inter-trial distances, however, decreased below baseline noise during the transient response phases (figure 4.4C), indicating that the trajectories during the dynamic phases of PN activity are less variable than either baseline or the odor-specific fixed point. Taken together, these data indicate that stimulus

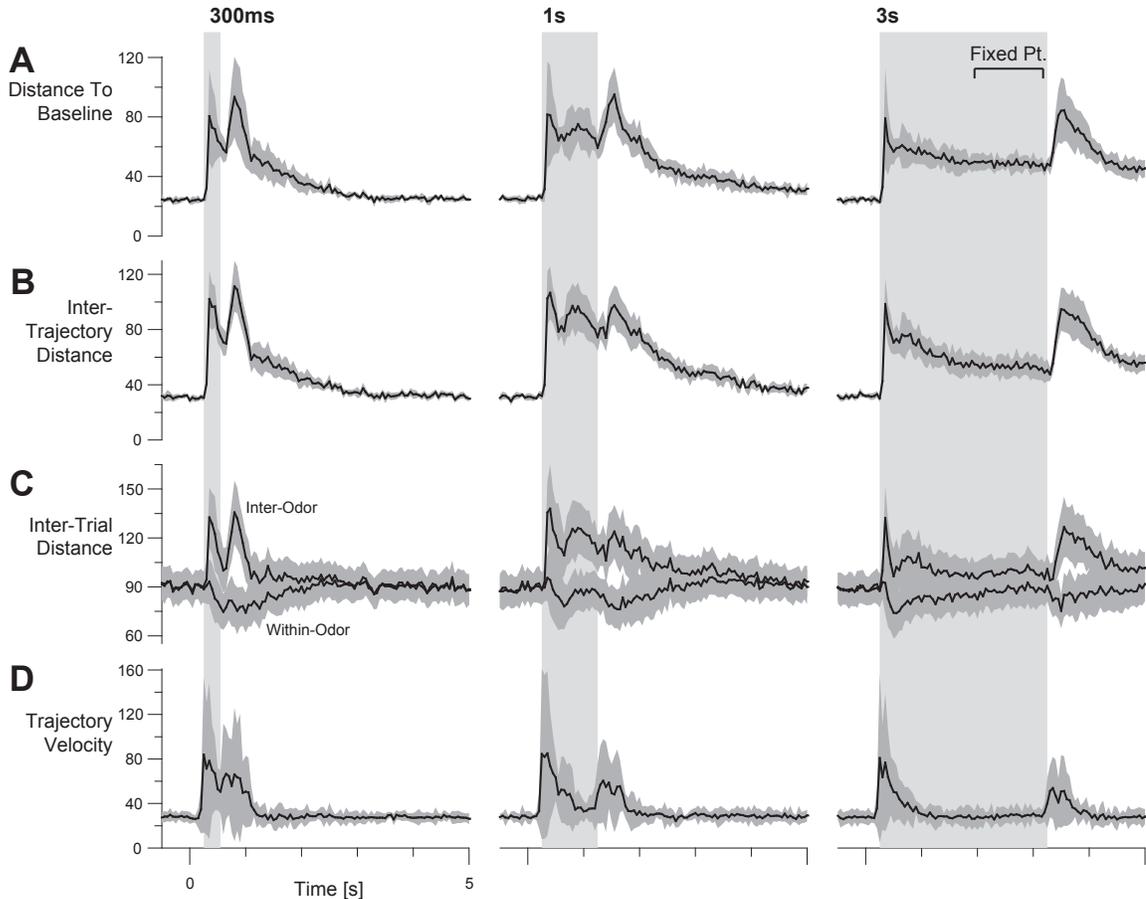


Figure 4.4. Quantification of inter-trajectory distances. All distances are computed as Euclidian distance in 99-dimensional space. *A*, instantaneous distance of mean odor trajectory to mean baseline state (average over five odors). Baseline state is computed from a 4 s average of pre-odor activity. *B*, instantaneous distances between mean odor trajectories (mean computed pair-wise across all five odors). *C*, mean instantaneous trial-to-trial distances, computed pair-wise across all trials from different odors (top traces) and within the same odor (bottom traces). Note that throughout the odor response, mean inter-trial distances across odors are substantially higher than within odor distances. *D*, instantaneous velocity of mean odor trajectories (average over five odors). All traces computed in 50 ms non-overlapping bins. Grey shading represents ± 1 SD for all subfigures. All plots in arbitrary units.

identification using instantaneous PN activity vectors is most accurate during the dynamic phases of the PN population response.

Using these measures of activity, we next estimated the rate at which the PN vectors change over time. For every 50 ms time bin, we calculated the instantaneous velocity of the ongoing trajectory by measuring the distance from that point to that 100 ms later. Figure 4.4D shows the mean and SD of this estimated velocity, sampled over the five odors. We found that the evolution of the PN vectors was faster at odor onset than at offset, and that steady state was reached within ~ 1.5 s of odor onset in the case of a sustained stimulus.

4.1.3 Information content of single trials

PN activity in response to multiple trials of one odor varied slightly from trial to trial (figures 4.1A, 4.3D). Our goal here is to assess the effect of this inter-trial variability on odor discrimination with instantaneous PN vectors. We thus measured, for each odor (cis, cit, hex, meth, and oct), pulse duration condition (0.3, 1, and 3 s), trial (2–10) and 50 ms time bin, the instantaneous PN vector characterizing the state of the 99 PNs. We show the PC projections of a subset of these vectors, calculated for four separate time bins: during baseline, 200 ms after odor onset, 2.5 s after odor onset (only for 3 s long odor pulses) and 200 ms after odor offset (figure 4.5A). We then used a simple algorithm (see methods) to classify each PN vector on the basis of its distance to the cluster centers of all five odors (in the corresponding time bin). One can see, for example, that the cis-3-hexen-1-ol (red) and 1-hexanol (cyan)

vectors overlapped during the fixed-point period (third panel, figure 4.5A) but were well separated at odor onset and offset (second and fourth panels, figure 4.5A). Interestingly, these two odors are chemically similar.

The percentage of correctly classified instantaneous PN vectors (over 45 individual trials: 9 trials per odor \times 5 odors) is plotted in figure 4.5B, time bin by time bin, for each one of the three odor durations. The dashed lines indicate chance level (20%, given 5 odors). For all three odor pulse durations, the percentage of correct classification was at chance level at baseline, rose to almost 100% within 100–200 ms after onset and decayed back to chance some time after odor offset at rates correlated with the duration of the preceding pulse (the shorter the pulse, the faster the return to chance-level classification). Correct classification remained close to 100% throughout the 0.3 and 1 s odor pulses, and at the beginning of the dynamic phase following odor offset). Correct classification rates fell to \sim 80% towards the middle of the 3 s responses. This corresponds to the times when odor-specific fixed points are reached, and when distances between odor trajectories are reduced (figure 4.4B).

We then used this classification method to test the temporal sensitivity of the decoding rule: specifically, we examined how the rate of successful classification of PN activity vectors (measured in one time bin) degrades as the classification templates are taken from time bins further and further away from that of the test vector. Figure 4.5C represents classification success as a function of time (x-axis) and of temporal offset between test and template vectors (y-axis). Offsets between

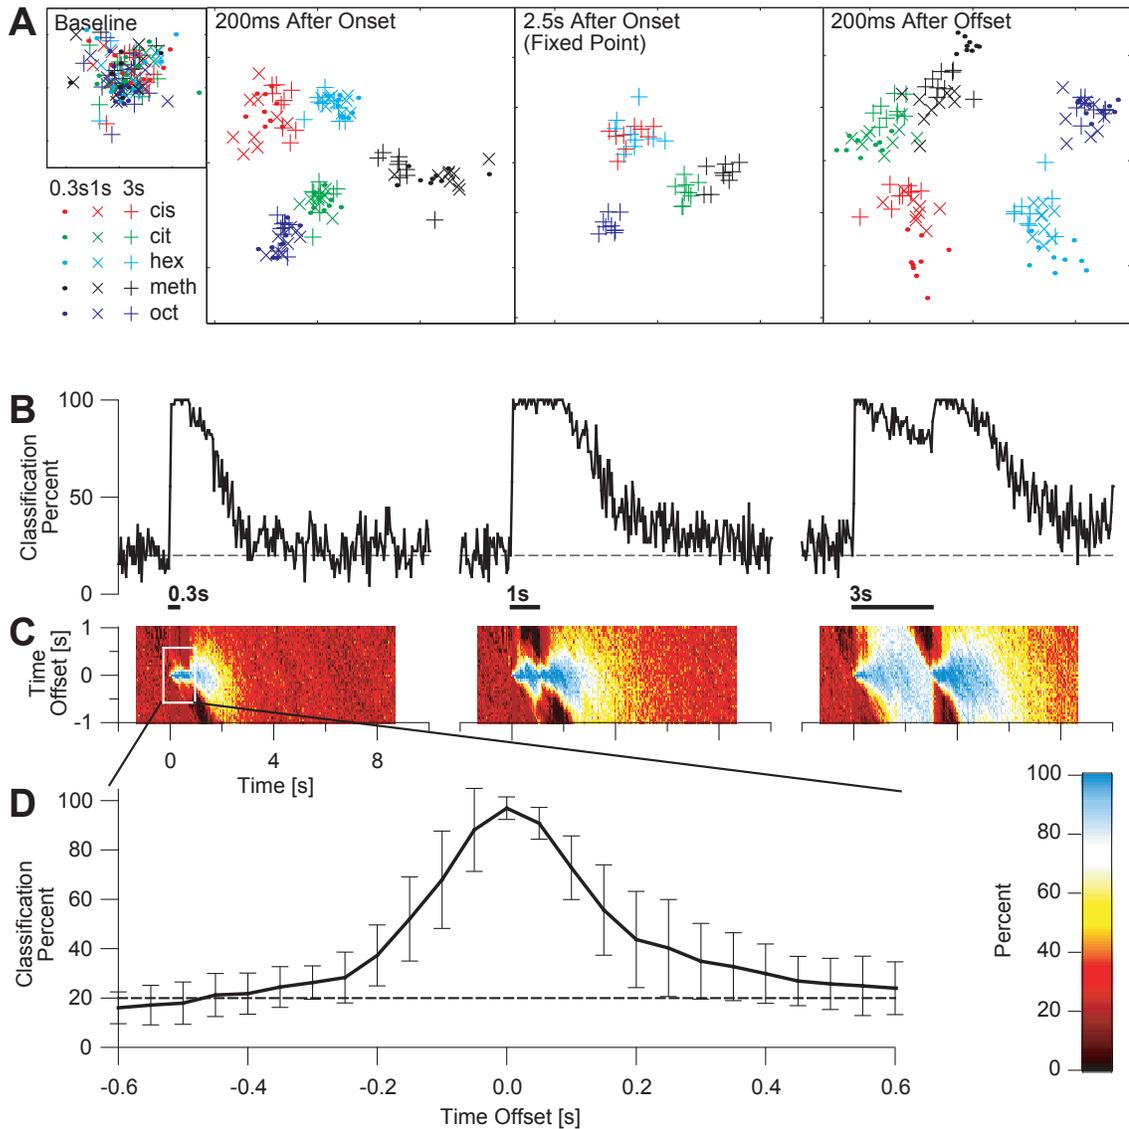


Figure 4.5. Odor information is available from short time slices of single trials. *A*, the activity of the PN population captured at a four different time points: baseline (pre-odor), 200 ms after odor onset, 2.5 s into the response (3 s odor only), and 200 ms after odor offset. The first two principle components are displayed, computed separately for each time period. Each marker represents the PN population activity (mean firing rate in 100 ms) from a single trial of one odor. Note that each odor forms a separate, non-overlapping cluster, and trials from different odor durations of the same odor are intermingled within each odor cluster. *B*, odor classification percentage over time (based on single-trial activity in 50 ms time bins, see methods). *C*, odor classification percentage over time with temporal offset between test and training data. *D*, effects of training/test offset on classification percentage, averaged over first 1 s of response to 300 ms odor pulse.

–500 ms and +500 ms were examined, in 50 ms increments. At zero offset ($y = 0$, figure 4.5C), classification success corresponds to the plots in figure 4.5B. The thickness or width of the blue band (see box, figure 4.5C) thus represents the temporal sensitivity of the decoding rule. This width was averaged over the first 1 s of the PN-ensemble response (300 ms odor pulse, 5 odors, 20 successive 50 ms bins), and plotted in figure 4.5D. Throughout this 1 s period, the average successful classification rate with no offset is 96.9% (SD: 4.52). With temporal offsets of ± 150 ms, however, successful classification rate dropped to 55.7% and 52.0% respectively (SDs: 18.3, 17.1). This shows how quickly the PN activity vectors evolve over time: a few oscillation cycles (50 ms long each on average) are therefore sufficient time to decorrelate PN-activity vectors with themselves enough to hinder odor identification using inter-vector distances.

4.1.4 *Local field potentials*

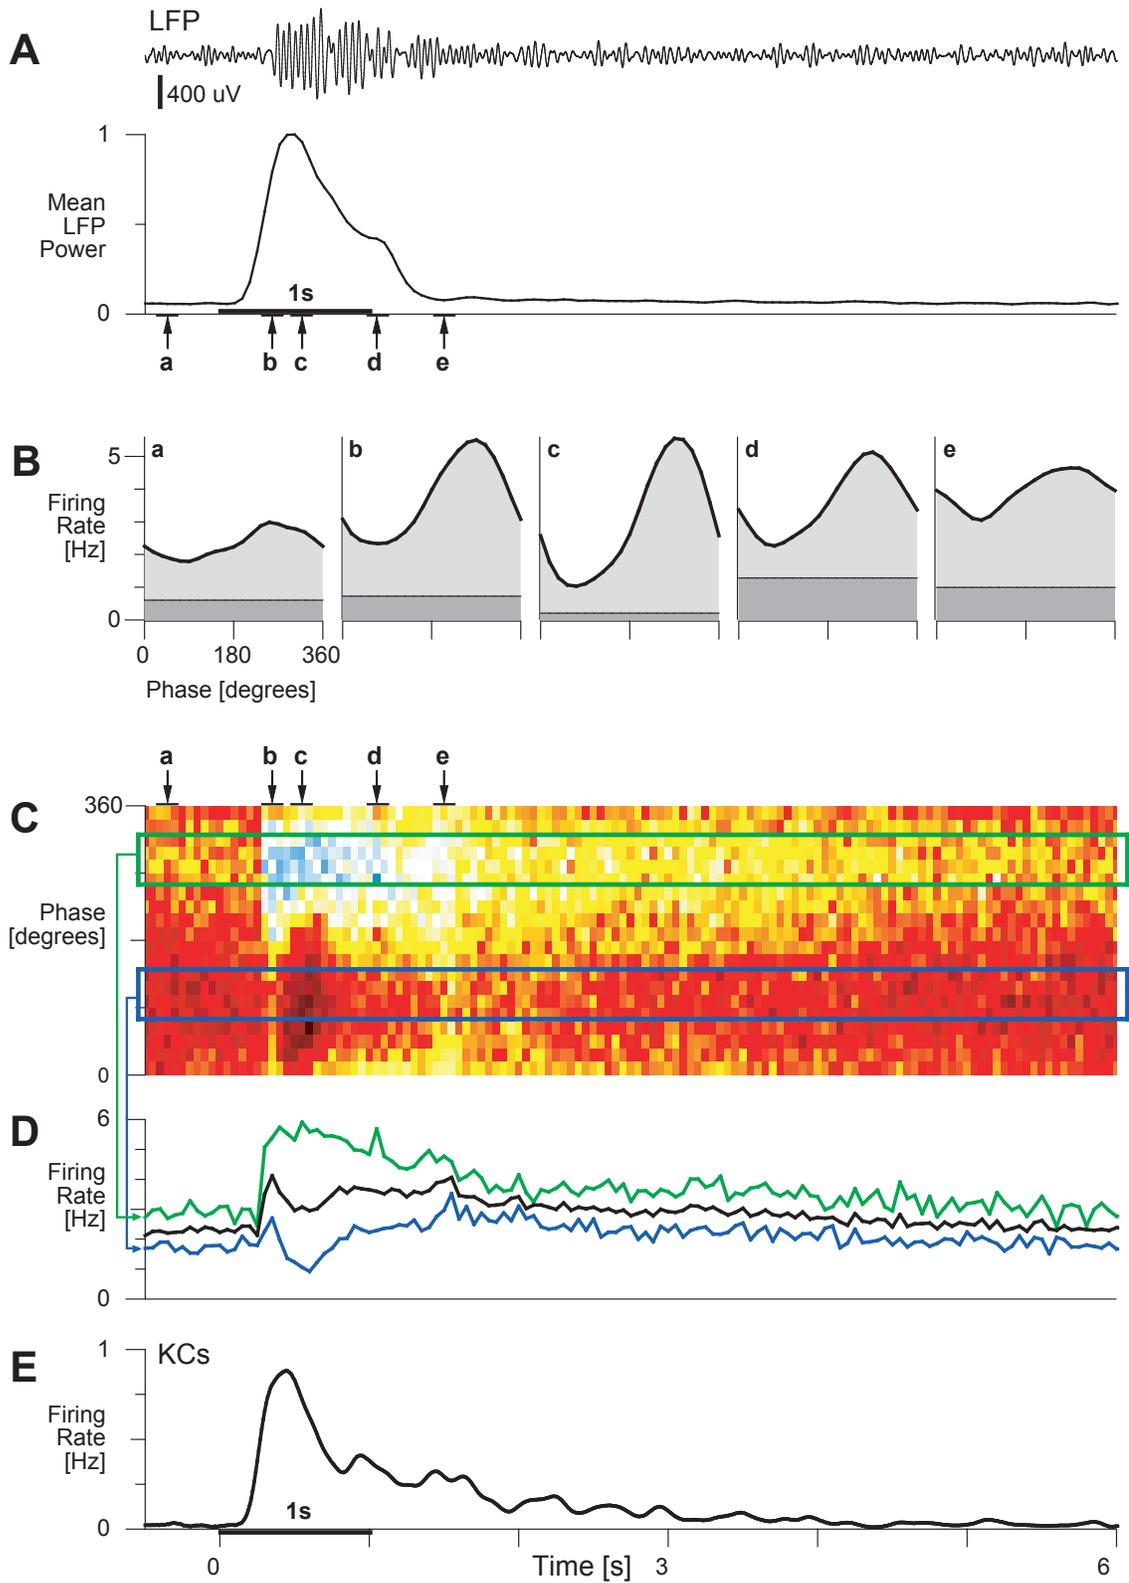
The dynamical PN response patterns examined above occur on top of faster periodic (20–30 Hz) synchronizing events across the PN populations (Laurent and Davidowitz, 1994; Wehr and Laurent, 1996). During an odor response, individual PNs are thus transiently synchronized with other PNs, in an odor-specific manner. Following our analysis of PN-response evolution over time (above), we study here how synchronization across the PN population evolves during an odor response.

LFPs are typically recorded with extracellular electrodes from the Kenyon cell body cluster, dorsal to the mushroom body calyx. These LFPs result mainly from the

alternating excitatory and inhibitory synaptic currents generated in Kenyon cells (the targets of PNs in the mushroom body) by PNs and lateral horn interneurons, respectively (Perez-Orive et al., 2002).

We measured the average LFP power (in the 10–30 Hz band) as a function of time around an odor pulse. Power was calculated over a scrolling window (width 200 ms, step 50 ms) over all trials (10 per stimulus condition) and odors (5), and averaged over 10 antennal lobes. This average is plotted for a 1 s odor pulse, together with a typical single-trial LFP trace (figure 4.6A). The single trial illustrates that short bouts of synchronization occur throughout baseline but that the dominant power emerges right after odor pulse onset. Mean oscillatory power reached a peak within about 300 ms of PN-population activation onset and decreased in two successive phases: the first occurred during the odor pulse; the second, after odor offset. The period of peak oscillatory power thus corresponds to the dynamic phase when

Figure 4.6 (on the next page). *A*, example single-trial LFP response (upper trace, bp filtered 10–30 Hz), and mean normalized LFP power in 15–25 Hz band (lower trace, mean over 5 odors and 10 recording sessions). Both traces for 1 s odor duration. *B*, smoothed phase alignment histograms for PN spikes with respect to the LFP, computed at 5 different time points during the response. The phase distributions are all shifted vertically (dark grey boxes) by an area corresponding to the number of spikes for which phase could not be attributed (due to insufficient 20 Hz power). *C*, 2-dimensional histogram of PN spike phase over time in response to a 1 s odor pulse. Each column represents a 50 ms time bin. The values along the column are the histogram of all PN phases (with respect to the LFP) recorded during that time bin. *D*, comparison of response strength for PN spikes of different phase alignments. Green trace (*top*): mean firing rate for spikes locked to the preferred phase of the LFP. Blue trace (*bottom*): mean firing rate for spikes locked to the non-preferred phase of the LFP. Black trace (*center*): mean firing rate for all spikes. Spikes with phases in the quarter cycle centered on 270° (90°) were considered locked to the preferred (non-preferred) phase. Firing rate for phase-aligned spikes represent probability of spiking within the preferred (non-preferred) quarter-cycle. *E*, mean Kenyon Cell firing rate in response to a 1 s odor stimulus (N=20 KCs, 6 odors).



the PN population vector evolves away from baseline, but before it reaches its odor-specific fixed point (figures 4.3 and 4.4).

Because all the PNs were recorded simultaneously with an LFP, we could attribute a phase to all the PN spikes that occurred when 20 Hz oscillatory power was significant. Figure 4.6B shows the distribution of PN spike phases, calculated at different epochs before, during and after the odor pulse. The phase distributions are all shifted vertically (dark grey boxes) by an area corresponding to the number of spikes for which phase could not be attributed (due to insufficient 20 Hz power). At rest, PN spikes showed a small degree of phase locking to the rising phase ($\sim 270^\circ$) of the LFP oscillations. (Such measures could only be made piecewise, with spikes produced during short spontaneous bouts of LFP oscillations, themselves likely due to uncontrolled odors activating the antenna.) During the odor response itself, PNs were strongly phase-locked to the rising phase of the LFP oscillations (mean: $\sim 290^\circ$). The preferred phase of PN firing relative to the LFP remained constant. By computing similar phase histograms for every time bin and odor trial, we constructed the 2-dimensional histogram in figure 4.6C. This histogram represents the relative number of spikes attributed to each phase (y: 18° steps) over time (x: 50 ms steps). It shows that, in the 1 s following odor onset, the total number of spikes with phases around 270° (rising phase) increased sharply. Furthermore, there was a transient decrease in spikes with phases between 0° and 180° (falling phase) during the first 300 ms of the response. We next considered only those spikes with phases within 36° of the mean preferred phase ($\sim 290^\circ$) and

measured their instantaneous rate of occurrence over time. Similarly, we measured the occurrence rate of all spikes falling within 36° of the mean non-preferred phase ($\sim 90^\circ$). Both measures are plotted in figure 4.6D (green and blue), along with the mean PN firing rate (black). The interval between these two bounds represents the proportion of well-locked PN spikes; this measure correlates better with the mean LFP-power profile (figure 4.6A) than does the mean PN firing rate. The increase in LFP power is thus more a result of increased correlation between PN spikes than of mean increased PN output.

4.1.5 *Kenyon cell responses*

We examined whether the detailed statistics of the PN population output could help identify those features of the PN response that are most likely to drive a Kenyon cell (KC) to respond. We showed previously that KCs are tuned to detect synchronous inputs, thanks to a combination of intrinsic properties and delayed feed-forward inhibition that decrease the KCs' integration time constant (Perez-Orive et al., 2002, 2004). Feed-forward inhibition is phase-locked to PN output (thus to the LFP) and follows it with a half-cycle delay (Perez-Orive et al., 2002). We thus hypothesized that KC activity might correlate not only with mean PN firing rate, but possibly better with PN oscillatory coherence.

We thus recorded from KCs and compared their mean distribution of response probability (1 s odor pulses, figure 4.6E) with several potentially relevant statistics of the PN population response. We first considered the two-second period following

odor onset, which includes 85% of odor-evoked KC spikes. During this time, the mean PN firing rate increased by 44% over baseline. In contrast, the rate of phase-locked PN spikes (figure 4.6D, green trace) increased by 57% over baseline. We next defined PN phase-locking strength as the difference between the occurrence rates of locked and unlocked PN spikes (i.e., the difference between the green and blue traces in figure 4.6D). This quantity increased by 111% over baseline during the same period.

To compare the time courses of these measures to the distribution of KC response probabilities, we computed the correlation coefficient (Pearson's coefficient, zero lag) between the mean KC firing rate (figure 4.6E) and each of the previous three PN statistics (in 50 ms time steps over 1 s following odor pulse onset). The correlations were 0.83 and 0.82 with phase-locked PN spike rate or PN phase-locking strength, respectively, but 0.42 with mean PN firing rate.

Finally, we tried to identify the epochs of the PN population response—onset dynamic phase, offset dynamics, or fixed point—that best evoke spiking in the KCs (those epochs should be the most relevant ones for olfactory processing). In response to 1 s odor pulses, KCs showed a detectable increase in firing over baseline for ~4 s (figure 4.6E); most KC spikes occurred over the 1 to 2 s after odor onset. To help separate the relevance of the three epochs of the PN population response, we also recorded the responses of KCs to 10 s long odor pulses. Across recordings from 20 Kenyon cells and 6 odors, the mean firing rate in the 1 s following odor onset was 0.49 spike/s., It was 0.24 spike/s in the 1 s following odor offset, and only 0.059

spike/s over the last eight seconds of the 10 s odor pulses. This rate, obtained during the odor-specific fixed point, was not substantially different from the baseline (pre-odor) rate of 0.045 spike/s, indicating that PN output at steady state—even though the fixed points are odor specific—produces little, if any, response across the KC population.

4.2 Discussion

4.2.1 *PN population responses to odors as trajectories in phase space*

The study of 99 locust antennal lobe projection neurons within a population of about 800 gives us the most precise picture so far of the collective behavior (firing probabilities and correlations) of this system during an odor stimulus. To facilitate its description, we measured the state of the PN ensemble as time series of PN activity vectors, measured over short, successive time bins. This strategy was critical to capture the evolution of this highly dynamical system. Our time bin was typically 50 ms, for reasons that will be justified later. The main results can be summarized as follows.

At baseline, individual PNs fire at an average rate of about 2.5 spikes/s. Their spikes are not significantly correlated with each other over short time scales (hundreds of milliseconds and less), except over very brief and randomly occurring epochs, corresponding to spontaneous encounters with ambient odor: such brief epochs of LFP oscillations disappear when the antenna is insulated or cut off. When represented in PN state space, baseline corresponds to a sphere of noise in

the neighborhood of a fixed origin (B, figure 4.7A), from which no information about stimulation history can be extracted. Shortly after odor pulse onset, activity spreads across the PN population, in a manner that depends on the duration of the stimulus. With long odor pulses (≥ 3 s), PN activity is characterized by three successive phases: (1) The first is an onset transient (on, figure 4.7A), when activity evolves rapidly across the PN population and is highly coherent; this phase can be represented geometrically as a stimulus-specific trajectory in PN state space. (2) After 1.5–2 s, an odor-specific fixed-point is reached (F, figure 4.7A); PN responses cease to evolve and coherence is low. Stability persists for at least 10 s if the stimulus is maintained. This phase is consistent with recent calcium imaging experiments in honeybees, suggesting the existence of odor-specific fixed points (Galán et al., 2004). (3) Following odor pulse termination, we observe a second transient (off, figure 4.7A), during which PN activity becomes dynamic again, with reduced coherence relative to the on-transient, and finally relaxes to baseline at decreasing velocity. The on- and off-transients are always different from one another yet odor specific.

With shorter odor pulses, the segments corresponding to an odor-specific fixed-point (or to the approach to it) are never experienced; the trajectories then consist of joined on- and off-transients, with diminishing excursions as the odor pulses become shorter. The off-transients following pulses of the same odor, but of different durations, generally converge before reaching baseline. No steady state is ever reached for odor pulses lasting 1 s or less; their representations across the

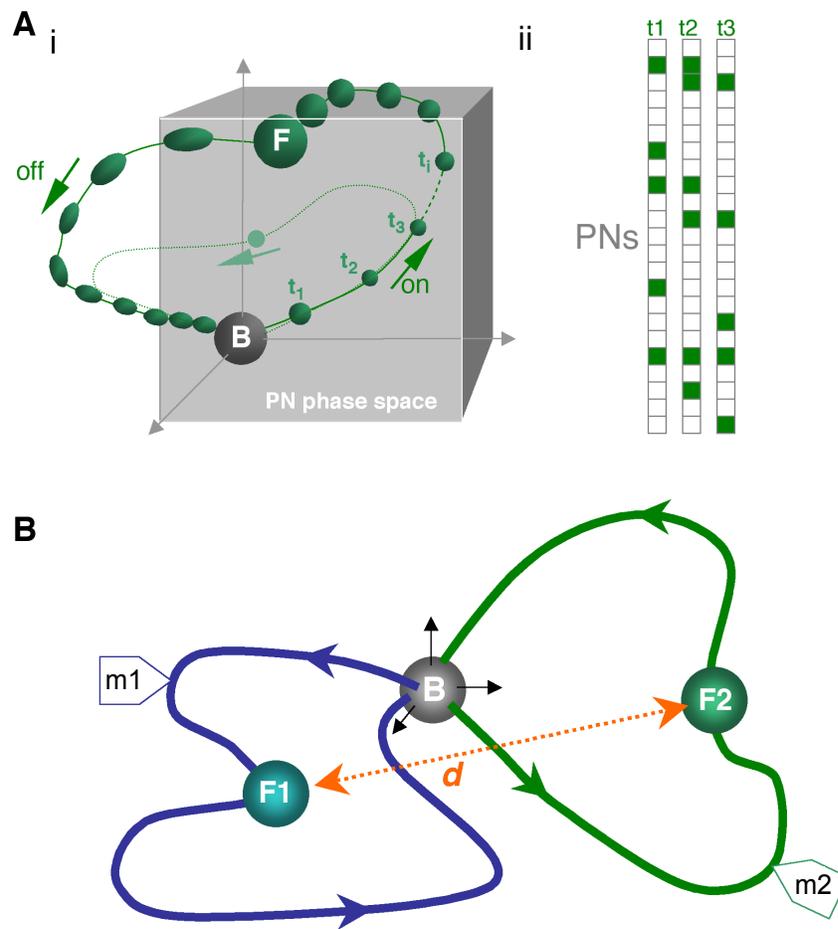


Figure 4.7. *Ai*, idealized odor trajectory in PN space. *Aii*, evolving PN activity underlying the trajectory in part (i). *B*, trajectories for two different odors. Trajectories differ at their fixed points (F1 and F2), but are maximally separate during the transient response phase (e.g., m1 and m2).

PN population are thus always non-stationary. In such conditions, it is debatable whether the distinction between on- and off-transients is still useful: practically speaking, the entire response period becomes one loop from and back to baseline, through a stimulus-specific sequence of stimulus-specific PN states. The shorter the stimulus, the shorter the loop. Trial-to-trial variability of the instantaneous PN activity vectors is less during either transient than it is at rest or at an odor-specific fixed point. We now examine the significance of these results.

4.2.2 Comparison to the winner-less competition framework

Recent theoretical work by Rabinovich and colleagues (Rabinovich et al., 2001; Laurent et al., 2001) proposed a nonlinear dynamical systems framework to describe odor-evoked PN activity. In this framework, called winner-less competition (WLC), PN activity can be thought of as stimulus-specific orbits in PN phase space: these trajectories are defined by sequences of unstable attractors (saddle states), that each corresponds to the activation of a specific subset of synchronized PNs. The attractors are the trajectories (heteroclinic orbits) themselves. They are very sensitive to the input, but globally stable: small individual-neuron deviations are corrected at the next step by the distributed and stabilizing effects of the population. When the input is withdrawn, the system relaxes back to the origin (baseline). This particular form of dynamical behavior can arise when connectivity between the network elements follows certain topological rules (e.g., asymmetrical connections, loops and particular statistics of connections among and between excitatory and inhibitory

neurons), and does not require perfect balance between excitation and inhibition. Whether the internal connectivity of the antennal lobe fulfills the appropriate topological constraints is not known. Our present results also revealed a deviation from the predictions of the initial WLC model: the existence of odor-specific fixed points for long odor pulse duration. More recent theoretical work, however, indicates that for at least one model network structure, conditions can be found such that a sequence of saddle points, linked by one-dimensional separatrices, lead towards a fixed-point attractor (Afraimovich et al., 2004). WLC dynamics can therefore be expressed transiently, at least in theory, on the path towards a fixed-point attractor. WLC is thus qualitatively still compatible with our experimental results. While WLC is the theoretical framework most consistent so far with our experimental observations, bridging the gap between experiments and theory remains a very difficult challenge.

4.2.3 How should one decode the PN trajectories?

Under many natural circumstances (short intermittent odor pulses, as typical in odor plumes), PN activity vectors have no opportunity to stabilize at odor-specific fixed-points. Average-rate measurements of PN responses are therefore inappropriate. An obvious question arises: if the PN vectors evolve, what is the appropriate time step over which the trajectories should be analyzed? To answer this, we turned to the PNs' own targets, the Kenyon cells, because they are the natural decoders of PN activity. Both their connectivity to the PN population and their integrative

properties helped us define the relevant spatio-temporal scales for analysis.

We chose an analysis window of 50 ms, because it captured the mean rate of collective PN output (20 Hz, as measured by the LFP or from the membrane potential of KCs during odor responses) while not confounding potentially faster evolution of PN output: indeed, we know from previous work (Wehr and Laurent, 1996) that PN spike phase (relative to the global 20 Hz output) does not vary with stimulus conditions. For example, we never observed phase precession, as observed in hippocampal place cells (O'Keefe and Recce, 1993; Skaggs et al., 1996), or odor/concentration-dependent PN spike phases (Stopfer et al., 2003), as proposed by Hopfield (1995): steps shorter than one oscillation cycle thus seemed to have no functional justification. While longer than the typical window over which synchronized PN output occurs ($\sim 10\text{--}20$ ms), the 50 ms window captured the fastest stimulus-linked evolution of the PN population output. At the same time, 50 ms was not so short a time as to increase noise unnecessarily in our estimates of instantaneous PN activity.

4.2.4 When is decoding of PN activity vectors most informative?

PN instantaneous firing rates averaged over the entire population increased by only ~ 1.5 spikes/s relative to baseline at the peak of the population response amplitude. What changed most was the variance of instantaneous firing rates across the population. The periods of peak variance were also those when vector classification was the easiest. Indeed, a result most surprising to us was that peak separation

of instantaneous PN activity vectors (averaged over trials) did not occur at the odor-dependent fixed points, but rather sometime during the dynamic phases of a response. Trial-to-trial variations were also smaller during the dynamic phases than at the fixed points. For example, two related alcohols could not be distinguished from their corresponding instantaneous PN vectors at the fixed point; they could be separated, however, during the on and off transients. These results indicate that the PN activity vectors (especially with single trials) are most informative during the dynamic phase immediately following PN population response onset. This period lasts no more than several hundred ms (estimated over this PN and stimulus set), but perfect discrimination was sometimes reached as early as within 100–200 ms. This transient period is characterized by peak PN synchronization, peak firing-rate variance across the PN population, fastest rate of PN-vector evolution, and lowest inter-trial variability.

We noted that the proportion of PN spikes that were reliable (i.e., those that occurred in the same 50 ms bin in at least 6 of 9 different trials with the same stimulus) during this optimal period was only about 60%. These spikes come from a population of about 80 PNs (10% of 800). Because ~480 PNs (60%) are silent at this point, the remaining 40% of spikes must come from a less-reliable population of ~240 PNs. This observation is interesting in the context of these circuits, for the following reason. The decoders of PN output (the Kenyon cells, KCs) are believed to be connected to a very large number of PNs (Jortner, Farivar, and Laurent, submitted). This means that a few KCs could receive input from all 80 PNs that

reliably respond to odor A at time bin i , as well as the most of the 240 less-reliable PNs. Such a KC would receive maximal input at time bin i during the response to odor A , regardless of which of the less-reliable PNs fired in a given trial. At all other time bins in response to odor A , and at all times during the responses to other odors, the input to this KC would be sub-maximal. This is because some of the reliable PNs at time i become unreliable or silent PNs at other times during the response. Hence, a KC can tolerate that a large proportion of its inputs be non-reliable, even at the time when its response is the most reliable. This was noted first (to our knowledge) by Kanerva (1988) in his mathematical exploration of sparse distributed memories in networks of binary units: distributed connectivity reduces overlaps, which provides great stability.

4.2.5 Comparison to Kenyon cell response profiles

Kenyon cells are the PNs' targets in the mushroom body. In locust, each Kenyon cell receives inputs from an estimated 400 PNs on average (Jortner, Farivar, and Laurent, submitted) and yet, hardly ever fires at rest and responds to very few odors (Perez-Orive et al., 2002). In the mushroom body, odors are thus represented by small assemblies of very specific neurons. This specificity was proposed to result from a combination of synaptic and cellular properties of KCs, feed-forward inhibitory circuits, connectivity and from the statistical structure of PN output (Perez-Orive et al., 2002, 2004). By measuring the instantaneous firing rate profiles of KCs over the same odors as used with PNs, we found here that KC response probability

correlates least with PN instantaneous firing rates, and best with instantaneous PN oscillatory coherence (in the 20 Hz band). Similarly, KC responses are least common when PN population activity reaches a fixed point during sustained odor stimuli. These results indicate that KC responses are driven primarily by the co-activation of synchronized PN assemblies, at precisely those phases of a response when the PN assemblies change the fastest. The antennal lobe code for odors thus consists of vectors of PN activity, updated once per oscillation cycle. Over a single trial, and during the response phase that corresponds to peak KC activity, PN activity evolves so rapidly that classification of PN activity vectors measured at one time against templates measured five cycles earlier or later is no better than chance (figure 4.5). This rapid evolution helps explain the temporal specificity we observed with individual KC responses (Perez-Orive et al., 2002, 2004): because KCs receive a strong voltage reset (IPSP) in the second half of each oscillation cycle (Perez-Orive et al., 2002; Laurent and Naraghi, 1994), temporal integration across successive cycles is limited or prevented. KCs thus each decode PN activity vectors one oscillation cycle at a time. The limited width of PN-vector correlation (the limited time over which PN activity vectors remain sufficiently self-similar) explains why KC response times should be narrowly distributed.

4.3 Methods

4.3.1 *Preparation and stimuli*

Results were obtained from locusts (*Schistocerca americana*) in a crowded, established colony. Young adults of either sex were immobilized, with one or two antennae intact for olfactory stimulation. The brain was exposed, desheathed and superfused with locust saline, as previously described (Laurent and Naraghi, 1994). Odors were delivered by injection of a controlled volume of odorized air within a constant stream of desiccated air. Teflon tubing was used at and downstream from the mixing point to prevent odor lingering and cross-contamination. Odors were used at 1% vapor pressure (PNs) or 10% vapor pressure (KCs), further diluted in the desiccated air stream. Ten trials of each odor were delivered for each of three different odor durations (300 ms, 1 s, and 3 s). The first trial of data was excluded from analysis to minimize effects of short-term response plasticity (Stopfer and Laurent, 1999). For each odor, all trials of 300 ms and 1 s duration were presented interleaved (each 300 ms trial was followed by a 1 s trial, and vice versa), followed by all ten trials of 3 s duration.

Odors presented and their abbreviations are as follows: cis-3-hexen-1-ol (cis), citral (cit), 1-hexanol (hex), methyl salicylate (meth), and 1-octanol (oct).

4.3.2 *Electrophysiology*

Extracellular recordings were performed with silicon probes which were generously provided by the University of Michigan Center for Neural Communica-

tion Technology (<http://www.engin.umich.edu/facility/cnct/>). Custom-built 16-channel preamplifiers and amplifiers were used for the recordings. Two tetrodes were used simultaneously. The preamplifier has a unitary gain, and the amplifier gain was set to 10,000 \times . Data from the four channels of each tetrode were filtered (custom-built amplifiers, band-pass filter: 0.3–6 kHz (for PNs and KCs); 1–300 Hz (for LFP); and 1 Hz–6 kHz (for simultaneous KC and LFP channels)), continuously acquired (15 kHz/channel, 12 bit), and stored to disk. For PN recordings, tetrodes were placed within the antennal lobe soma clusters, peripheral to the neuropils at depths less than 200 μ m. During MB recordings (KCs, LFP), probes were either pressed on the surface of the MB or placed within the KC soma clusters, peripheral to the neuropils at depths less than 200 μ m. Cell identification was unambiguous because PNs are the only spiking neurons in the locust antennal lobe, (LNs do not produce sodium action potentials), and because all the somata located above the MB calyx belong to KCs.

Tetrode recordings were analyzed as described in chapter 2 and Pouzat et al. (2002). Briefly, events were detected on all channels as voltage peaks above a preset threshold (usually 2.5–3.5 times each channel's SD). For any detected event on any channel, the same 3 ms window (each containing 45 samples) centered on that peak was extracted from each one of the four channels in a tetrode. Each event was then represented as a 180-dimensional vector (4×45 samples). Noise properties for the recording were estimated from all the recording segments between detected events, by computing the auto- and cross-correlations of all four channels. A noise

covariance matrix was computed and used for noise whitening. Events were then clustered using a modification of the expectation maximization algorithm. Because of noise whitening, clusters consisting of, and only of, all the spikes from a single source should form a Gaussian ($SD = 1$) distribution in 180-dimensional space. This property enabled us to perform several statistical tests to select only units that met rigorous quantitative criteria of isolation.

A total of 99 PNs were recorded from 10 antennal lobes in 9 different locusts.

4.3.3 *Data analysis*

All data analysis was performed using Igor Pro (WaveMetrics, Lake Oswego, OR) or Matlab (The MathWorks, Natick, MA). Our response metric was computed by taking spike time data from each PN, odor, and duration and dividing data into successive, non-overlapping 50 ms bins. For each bin, the number of trials with at least one spike was computed. A PN was considered active if this count reached at least six (out of nine total trials).

Population trajectories were computed for each odor and stimulus duration as a series of vectors representing the state of the system over the course of an odor delivery trial. The vector representing the instantaneous state of the system was computed in successive, non-overlapping 50 ms bins. In each time bin, a 99-dimensional vector was generated in which value of the N th dimension was the mean firing rate of the N th PN. Mean firing rates were computed by taking the mean of the spike count in the 50 ms bin across all 9 trials of an odor presentation.

Odor trials were classified separately for each odor duration, based on the population activity vectors calculated on 50 ms bins. In each time bin, the centroid of activity for each of the five odors was computed by calculating the mean activity vector across the nine trials of each odor. Next, for each trial and each odor, the distance from the single trial population activity to each centroid was computed. Individual trials were attributed to the odor with the nearest centroid. The trial in question was always excluded from the calculation of centroids.

4.3.4 Local field potential and spike phase

Raw LFP signal from MB tetrode recordings (see section 4.3.2) was bandpass filtered from 10–30 Hz (non phase-distorting Butterworth filter, built-in Matlab functions). To compute spike phase histograms, each PN spike was assigned a phase with respect to the simultaneously recorded LFP using a simple algorithm. First, all peaks and troughs that exceeded 0.5 SDs of baseline (non-odor evoked) fluctuations were detected in the bandpass filtered LFP. Spikes were assigned a phase if they fell between a peak and a trough (in either order), and if less than 35 ms separated the peak and trough. Phase was then assigned based on a linear scaling of phase values between a peak and trough (0° – 180°), or trough and peak (180° – 360°). Spikes with no phase attribution were distributed uniformly across all phases.

CHAPTER 5

Concluding Remarks

THE PREVIOUS THREE CHAPTERS presented the significant results of the research projects I participated in while in graduate school. In the present chapter, I will summarize the main findings, briefly describe their significance, and highlight a few open questions raised by these results.

5.1 Spike-Sorting with Quality Control

5.1.1 *Summary of results*

The work presented in chapter 2 describes a novel, automated method for extracting the spike trains of simultaneously recorded neurons from raw multi-channel extracellular data (i.e., spike-sorting). The results show that a simple model of data generation combined with an accurate noise model can explain the raw signals recorded from the locust antennal lobe. Using this model, spike-sorting can be performed quickly and automatically using a variant of the EM algorithm (Dempster et al., 1977). The model makes several quantitative predictions about the statistics

of the spike-sorting results, which can be tested using a series of three quantitative tests described in section 2.1.4. The tests are shown to detect several common conditions that can generate unreliable sorting results (Lewicki, 1998). Furthermore, these tests are modular and can be adapted to other spike-sorting algorithms.

5.1.2 *Significance of results*

Spike-sorting is an extremely challenging problem to which many solutions have been proposed (Lewicki, 1998). Nevertheless, there is no universally accepted method that works well in all cases. In fact, it is often difficult to determine if a spike-sorting method works correctly or not. Most attempts to assess the quality of spike-sorting algorithms validate the spike-sorting results using simultaneously recorded intracellular data (Wehr et al., 1999; Harris et al., 2000). In contrast, the fundamental advantage of the method proposed in chapter 2 is that the result of the sorting can be validated based solely on the statistical properties of the extracellular recording. In most experimental paradigms, intracellular recordings are difficult or impossible. Moreover, intracellular validation is only informative about one extracellular spike train at a time. With statistical validation, the quality of all putative spike trains can be tested.¹

A second distinct advantage of the algorithm described in chapter 2 is that it is automated. Manual control of spike clustering is a significant source of error and unreliability (Harris et al., 2000). Manual sorting algorithms are also time

¹It should be noted that with intracellular validation, the true spike train is completely known, while statistical test will always operate within a margin of error.

consuming and do not scale well to large projects. Current methods allow for simultaneous recordings from over 50 channels (e.g., Siapas et al., 2005); sorting such quantities of data manually is unrealistic. Clearly, any sorting algorithm to be used for such large data sets must be fully automated, and must include an automated means of assessing result quality. The present algorithm meets these criteria.

5.1.3 Future directions

The sorting algorithm described in chapter 2 is based on a simple model of spike generation. Many cell types in the nervous system are known to have more complex behaviors. For example, the spike waveforms of cortical neurons often change during the course of a spike burst. While the present model cannot account for such non-stationarities, the general framework, including the quality tests, could be adapted to a more complex model, especially if the non-stationarities have recognizable structure. One such modification has already been developed by Pouzat et al. (2004).

An equally important direction for future improvement is the ability of the algorithm to handle long stretches of data. During an extracellular recording, the shape of the spike waveforms from the same cell often drifts over the course of many minutes. In some recordings this drift can cause two previously distinguishable cells to become statistically inseparable. Currently, the algorithm can be adapted to track the slow drift of distinguishable clusters (unpublished data). Nevertheless,

there is a strong need to extend the quantitative tests to explicitly include timing information and to incorporate waveform drift as well as the loss/gain of waveforms over time into the model.

5.2 Population Coding in the Locust Olfactory System

5.2.1 *Summary of results*

In chapters 3 and 4, the methods from chapter 2 are used to study how olfactory information is represented in two populations of neurons in the locust brain. Odors are represented by the distributed activity across the PN population. Kenyon cells (KCs) receive direct excitatory PN input; yet they encode odor information very differently. KC odor responses are significantly sparser than PN responses (as measured using lifetime and population sparseness metrics), and are much briefer when they do occur: the average KC response consists of fewer than 3 spikes, while an average PN response contains ~ 20 . This transformation is caused by a number of underlying factors working in concert.

- During an odor response, PN spikes tend to be phase locked to the odor-evoked ~ 20 Hz LFP oscillation—and therefore synchronized to other PN spikes. Kenyon cells, in turn, are tuned to synchronous input. Active membrane conductances within the KC and delayed feed-forward inhibition mediated by LHIs combine to decrease the KCs' integration time constant.
- Kenyon cells odor responses are characterized by regular, phase-locked IPSPs

that tend to occur $\sim 180^\circ$ out of phase with PN spiking. This periodic inhibition acts to reset the KC membrane potential, effectively causing the KC to “forget” any excitatory input received in the previous LFP cycle. Kenyon cells therefore fire (or do not fire) an action potential based primarily on their synchronous excitatory input during single cycles of the LFP, each cycle lasting approximately 50 ms.²

- During each cycle of the LFP, the excitatory input to a KC is determined by the set of PNs that connect to it and activity of those PNs at that time. The PN odor response evolves dynamically such that only a relatively small ensemble of PNs are co-active at a single oscillation cycle within the response. Therefore, a KC will only respond during those cycles when a significant fraction of its PN subset is active.

A better understanding of the process underlying the KC responses requires detailed knowledge about the dynamics of the PN population response. The work presented in chapter 4 begins to address this need.

The KC population reads out the PN population activity at a timescale of a single oscillation cycle (~ 50 ms) or less. Thus, this is the appropriate resolution with which to measure the statistics of the PN population response. In an average 50 ms time bin during an odor response, approximately 10% of PNs show a reliable response. In contrast, over 50% of PNs are reliably silenced within the same time

²It remains to be tested whether there is significant integration beyond single oscillation cycles. For example, short term synaptic plasticity at the PN-KC synapse could be such that successive oscillation cycles are not completely independent.

bin. The ensemble of responsive PNs changes over time in an odor-specific manner. Although there is some trial-to-trial variability in the odor response, the single trial PN population activity from only one 50 ms window contains significant information about odor identity.

These epochs of odor-specific dynamics do not last indefinitely. Instead there are three separate phases to an odor response. There are two distinct responses characterized by periods of strong odor-specific dynamics, one to the onset of an odor pulse, and one to the offset. Additionally, in response to odor durations longer than ~ 2 s, the PN population reaches a state of constant activity. While still somewhat odor-specific, this period is substantially less informative about odor identity than either of the two transient response phases.

5.2.2 Significance of results

These results establish the locust olfactory system as a model system for studying both the encoding and decoding of dynamical population codes. Previous studies have already described the odor-evoked spatio-temporal dynamics in the locust AL (Laurent and Davidowitz, 1994; Wehr and Laurent, 1996; MacLeod and Laurent, 1996). The work presented here further quantifies the PN population response and then examines how that signal is transformed into KC responses.

The individual components underlying this transformation are not unique to this system. For example, coincidence detectors have been observed in the auditory system. Carr and Konishi (1990) showed that neurons in the brain stem of the barn

owl work as precise coincidence detectors in order to discriminate interaural time differences.

Other components have been observed elsewhere but their precise roles in neural processing remain unclear. For example, oscillatory synchrony has been observed in many different systems (Adrian, 1942; Gelperin and Tank, 1990; Fries et al., 2000; Kreiter and Singer, 1996; Maldonado et al., 2000; Buzsaki and Draguhn, 2004), yet its precise role (if any) in many of these systems is hotly debated. Synchrony has been proposed to underlie perceptual binding (Singer, 1999) and attention (Fries et al., 2000), among other processes. These proposals have yet to be definitively evaluated.

In contrast, the results presented in chapters 3 and 4 provide a detailed example of a system where oscillatory synchrony, coincidence detection, feed-forward inhibition, and temporal patterning operate in concert to transform a dense population code into a sparse code. The existence of these same components in other systems support the idea that the concepts derived from studying the insect brain may transfer to other systems.

The results also firmly establish the existence of a temporal code (see section 1.1.2, p. 6) in the locust AL and quantify the speed and duration of its dynamics. Furthermore, these results demonstrate that even when an odor stimulus is very short (on the same timescale as the dynamics themselves), temporal dynamics play a role in the representation of that stimulus. A simple 300 ms odor pulse evokes dynamics that last for nearly a second.

5.2.3 *Open questions*

Projection neuron population responses can be represented as trajectories in PN phase space. In chapter 4, the trajectories in response to five different odors were examined. This represents only a small sample of all possible odor trajectories. One exciting future direction is to explore the set of response trajectories to a much larger set of odors and ask whether these trajectories themselves can be classified or ordered in informative ways. For example, do similar odors have partially overlapping trajectories (Friedrich and Laurent, 2001)? What is the frequency of overlap for non-similar odors? Does valence (whether an odor is attractive or aversive) constrain the trajectory of an odor? Is there a fundamental difference between the onset and offset trajectories or the information they encode? Finding and describing such patterns may be especially informative given what is now known about how KCs read the PN population output.

Another open question is how the KC responses are read out by the output neurons of the MB, the α - and β -lobe cells. Understanding this process would likely yield new understanding about the way KCs encode information and may provide further explanation for the sparse nature of their responses. For example, we do not know if the relative timing of KC spikes have a strong influence on the response of the α - and β -lobe cells.

Because odor representations in the KC population are so sparse, a lot of odor information is contained in the identity of the KCs that respond to an odor, without regard to their relative timing. It is possible that α - and β -lobe cells have long

integration time constants (e.g., hundreds of milliseconds), and are therefore less sensitive to the precise timing of KC spikes. Such a result would provide a potential explanation for the nature of the KC code: the sparse representation of odors may be sufficiently informative to allow timing information to be ignored, thus simplifying decoding.

APPENDIX A

Projection Neuron Baseline Statistics

In chapters 3 and 4, we characterized the response of the PN population to an odor stimulus. In this brief study, we characterize and model the spontaneous spiking activity of locust PNs in the absence of odor stimuli. We then test for correlations in the spontaneous activity of simultaneously recorded PNs.

A.1 Results

A.1.1 *Measuring firing rate statistics*

In order to characterize baseline PN activity, we made multi-unit extracellular recordings of PNs and collected long stretches (tens of minutes) of spontaneous spiking activity in the absence of odor. For each PN recorded, we began by computing its inter-spike interval (ISI) density (figure A.1B),

$$ISI(t) = p(ISI_i = t),$$

where $ISI_1 \dots ISI_n$ are all the inter-spike intervals recorded for that PN.

Following Johnson (1996), we next compute the hazard function from the ISI distribution (figure A.1C). The hazard function is defined as

$$\text{Hazard}(t) = \frac{ISI(t)}{\int_{\alpha=t}^{\alpha=+\infty} ISI(\alpha)d\alpha}, \quad (\text{A.1})$$

where $ISI(t)$ is the value of the ISI density for $ISI = t$. The hazard function is an intuitive way of understanding the behavior of a spiking cell. $\text{Hazard}(t)$ represents the instantaneous firing probability of a neuron, given that its previous spike occurred t seconds ago.

STATIONARITY TEST

In order for our computed statistics (the ISI density and hazard function) to be meaningful, we need to make sure the firing statistics of the PNs do not change significantly during the time we compute these statistics. In other words, we must make sure their firing statistics are stationary. We test for stationarity using a simple test. For each PN, all ISI values are binned into non-overlapping groups of 100 successive ISIs. The mean of each group of 100 is computed and plotted as in figure A.2A. If the firing statistics for this PN are stationary, the mean ISI value should remain constant over time, although with some amount of noise. We can derive a predicted value of the mean ISI, as well as a 95% confidence interval (for the 100 ISI averages) from the mean and SD of the entire set of ISI values (Brandt, 1999). The expected mean and 95% confidence interval are also plotted in figure A.2A.

The ISI data plotted in figure A.2A are typical for a locust PN. The measured

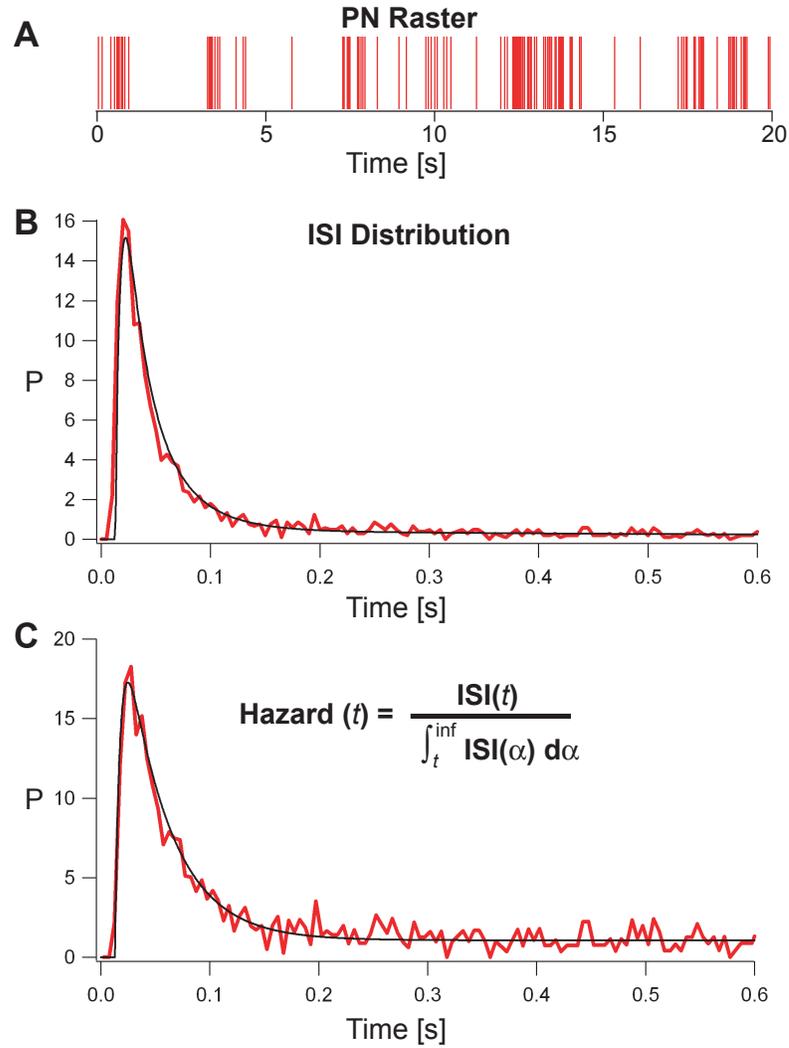


Figure A.1. Example of ISI distribution and hazard function. *A*, a 20 s sample of baseline spiking activity from one PN. *B*, the distribution (density) of inter-spike intervals (ISIs) for the same PN (*red*). ISI distribution was computed from 15 minutes of spontaneous activity with dry, unodorized air flowing across the antenna. *C*, hazard function (*red*) for the same PN. The hazard function is calculated directly from the ISI distribution as shown. *B* and *C*, black trace represents the estimated ISI density and hazard function after optimization of model parameters (see section A.1.2).

statistic does exceed the confidence interval more than 5% of the time (18% for the plotted data). Nevertheless, these are no strong trends that deviate from the expected mean over time. In a typical recording with tens of minutes of data from 5–10 PNs, there are often periods of time where not all PNs are stationary. Care was taken to only use data from a subset of PNs and a subset of the full recording duration, such that all PNs analyzed were relatively stationary for the entire duration of analysis.

ISI INDEPENDENCE TEST

A second fundamental question we asked about the statistics of PN spike trains is whether the timing of a PN spike is dependent on past history of the PN's spiking activity. Specifically, we asked if the timing of a PN spike depends only on the timing of the previous spike. This property would make PN spike trains analogous to the mathematical concept of a *renewal point process* (for more details see Johnson, 1996). Renewal processes have been studied at length, and are typically easier to study and model than other types of point processes.

If, in fact, the timing of a PN spike depends only on the timing of the previous spike, then the duration of successive ISIs should be statistically independent. In order to test for the independence of successive ISIs, we consider all pairs of successive ISIs. We then group the pairs based on the first ISI value in the pair, using non-overlapping 10 ms bins. For each 10 ms range of first-ISI values, we compute the mean second-ISI value. In this way, we compute the mean ISI, conditional on

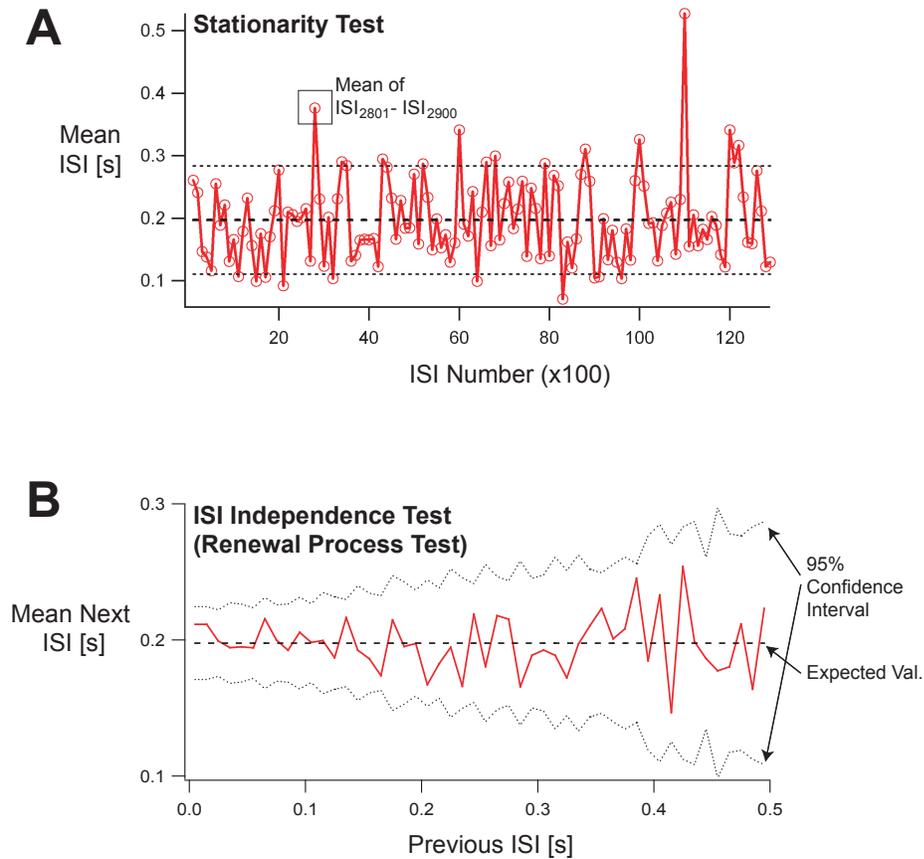


Figure A.2. Examples of two statistical tests for PN spike trains. *A*, a simple test for stationarity of a PN spike train. Non-overlapping groups of 100 successive ISIs are averaged and that series of values plotted in red. The mean over all ISIs is represented by the thick dotted black line. The 95% confidence intervals are plotted around the mean. For a stationary spike train, only 5% of the points should lie outside the confidence boundary. Additionally, slow trends in the data can be easily visualized with such plots. For the plotted data, 18% of the data points are outside the confidence interval, indicating that the the data is not purely stationary. Nevertheless, we observe no strong trends in the mean ISI. *B*, a test of ISI independence. In order to test for the independence of successive ISIs, we plot the mean ISI conditional on the previous ISI value (using 10 ms bins): $ConditionalMeanISI(x) = \langle ISI_n \cdot p(ISI_n | ISI_{n-1} = x) \rangle$. As in part *A*, the expected ISI (assuming independence) is plotted along with a 95% confidence interval.

the ISI which previously occurred:

$$\text{ConditionalMeanISI}(x) = \langle \text{ISI}_n \cdot p(\text{ISI}_n | \text{ISI}_{n-1} = x) \rangle,$$

where, in our case, x represents the first ISI of a pair, and takes on values in 10 ms increments. In figure A.2B we plot the *ConditionalMeanISI* function (*red*) for a simulated spike train generated from a Poisson model (Johnson, 1996). If successive ISIs are indeed independent, *ConditionalMeanISI*(x) should always return the same value—the mean ISI of the entire data set. Furthermore, we can also compute a 95% confidence interval around this expected value, based on the number of data points used to compute *ConditionalMeanISI*(x) for each value of x (Brandt, 1999) as shown in figure A.2B.

Figure A.2B shows the result of the ISI independence test when performed on a simulated Poisson spike train. As expected, this test shows that successive ISIs are independent, as they should be for a Poisson spike train. When the ISI independence test is run on real PN spike trains, most PNs showed a slight deviation from independence. Most PNs had a short ISI (< 100 ms) following another short ISI, though this trend was not always significant. These results will be further elaborated in the next section and in figure A.4.

A.1.2 Modeling PN baseline firing as a renewal process

Data collected from 35 PNs revealed that most PNs showed similar-looking ISI densities and hazard functions. We attempted to fit the hazard functions of all

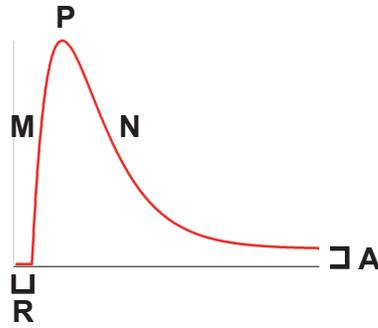


Figure A.3. Schematic of hazard model parameters. The hazard model is defined as: $HazardModel(t) = P(1 - e^{-M(t-R)})(\frac{A}{P} + e^{-N(t-R)})$ where R is the refractory period in seconds, M and N are time constants describing the rise and fall of the hazard function, A determines the final asymptotic level and P is a scaling factor.

PNs recorded to a simple model with only five parameters. The formula chosen for the hazard model allows for a refractory period (where $hazard(t) = 0$), an exponential rise in the hazard function, followed by an exponential fall which reaches an asymptotic value. The model is:

$$HazardModel(t) = P(1 - e^{-M(t-R)})(\frac{A}{P} + e^{-N(t-R)}),$$

where R is the refractory period in seconds, M and N are time constants describing the rise and fall of the hazard function, A determines the final asymptotic level and P is a scaling factor (see figure A.3).

The firing statistics for each PN were fit to this model by first computing the empirical hazard function for each PN (see equation A.1 and figure A.1B). Next, the five hazard model parameters were optimized to the PN's hazard function using Marquardt minimization (Brandt, 1999, p. 32). A model hazard function, as well as ISI density distribution, could then be computed for that PN.

Figure A.4 shows a comparison between empirical hazard functions and ISI distributions (*red traces*) and their models (*black*) for seven simultaneously recorded PNs. In addition, as a visual indication of goodness-of-fit, the model and empirical cumulative ISI distributions are also plotted. For comparison, the final two columns of figure A.4 show the same distributions computed for two Poisson spike trains (one with no refractory period, one with a 25 ms refractory period). These two Poisson distributions have a comparable firing rate (4 Hz) to the mean firing rate of all PNs recorded in this study. A total of 35 PNs were recorded and modeled and summary statistics for the 35 model hazard functions are plotted in figure A.5.

The two statistical tests described in section A.1.1 were also performed on all 35 PNs. The bottom two rows of figure A.4 show the results of running the two test on the seven PNs (and two Poisson models) of that figure. The stationarity test demonstrates that these seven PN are mostly stationary during this recording and show no strong trends towards increasing or decreasing their mean ISIs over time.

From figure A.4 we also see that the ISI independence test reveals a trend toward dependence between successive ISIs (where short ISIs are followed by short ISIs). We tested the statistical significance of this trend on a cell-by-cell basis. For each PN, we first computed the correlation coefficient of its entire series of successive-ISI pairs. We then shuffled the order of the PN's ISIs and recomputed the successive-ISI-pair correlation coefficient for 200 different shuffled configurations. The original correlation coefficient was considered significant if it fell outside the central 95% of the shuffled distribution. The p -values displayed under the ISI independence tests

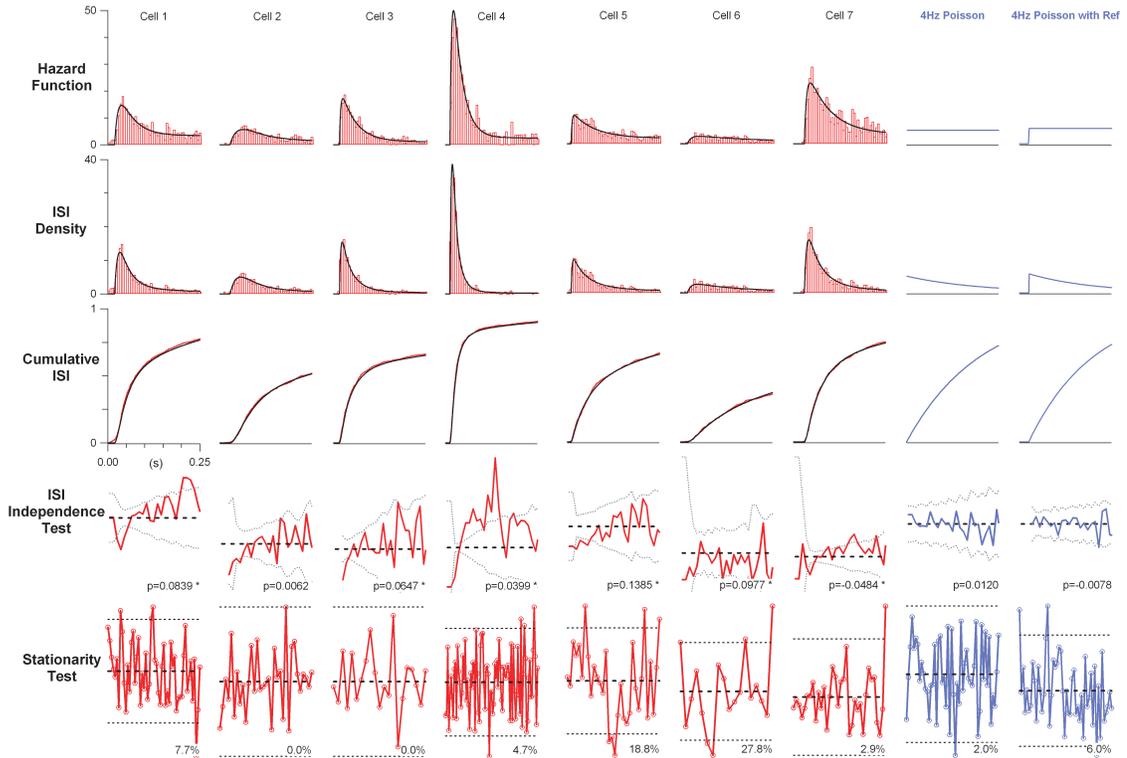


Figure A.4. Statistics from seven simultaneously recorded PNs. The first three rows show the raw (*red*) and modeled (*black*) versions of the hazard function, ISI density and cumulative ISI. Note the high degree of correspondence between the raw values and the model. The next row shows the result of the ISI independence test for each PN. This test shows a small, but consistent deviation from independence at short ISIs (a short ISI is likely to be followed by another short ISI). See text for an explanation of the p -values. The final row shows the result of the stationarity test. The fraction of points that exceed the 95% confidence interval is also displayed. None of the seven cells show any trend in mean ISI duration, although several cells do exhibit slightly higher variability than expected. The eighth and ninth columns show the same statistics computed on simulated data generated from the two simplest discharge statistics: Poisson and Poisson with a refractory period.

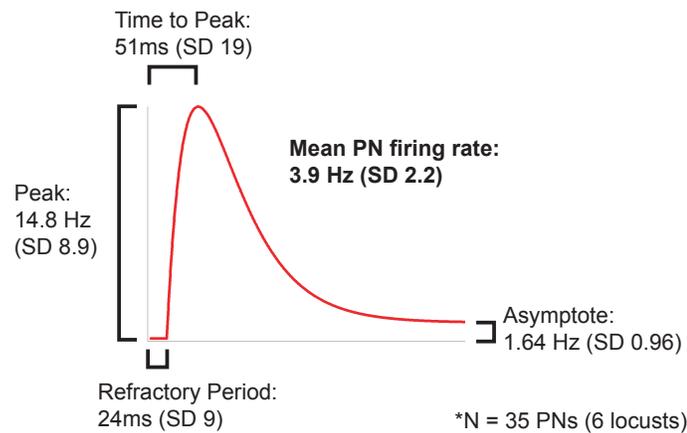


Figure A.5. Average hazard function statistics from 35 PN (6 locusts).

in figure A.4 show the result of this comparison.

The p -values displayed in figure A.4 indicate that for many PN successive ISIs are not purely independent. Therefore the statistics governing PN spike generation cannot be described entirely by a renewal process. Nevertheless, the dependence between successive ISIs is often small, and renewal processes are relatively straightforward models. For these reasons, we attempted to model PN firing with a renewal model.

In a renewal process, where successive ISIs are independent, the ISI distribution (or the hazard function) completely characterizes the the behavior of the model. To generate simulated data from a renewal process, one can produce a set of spike times where the interval between each spike and the next is sampled from the model's ISI distribution (Johnson, 1996). Using this technique we were able to generate simulated spike trains for each of the recorded PN using their model ISI distributions. A comparison between raw and simulated spike trains for the same

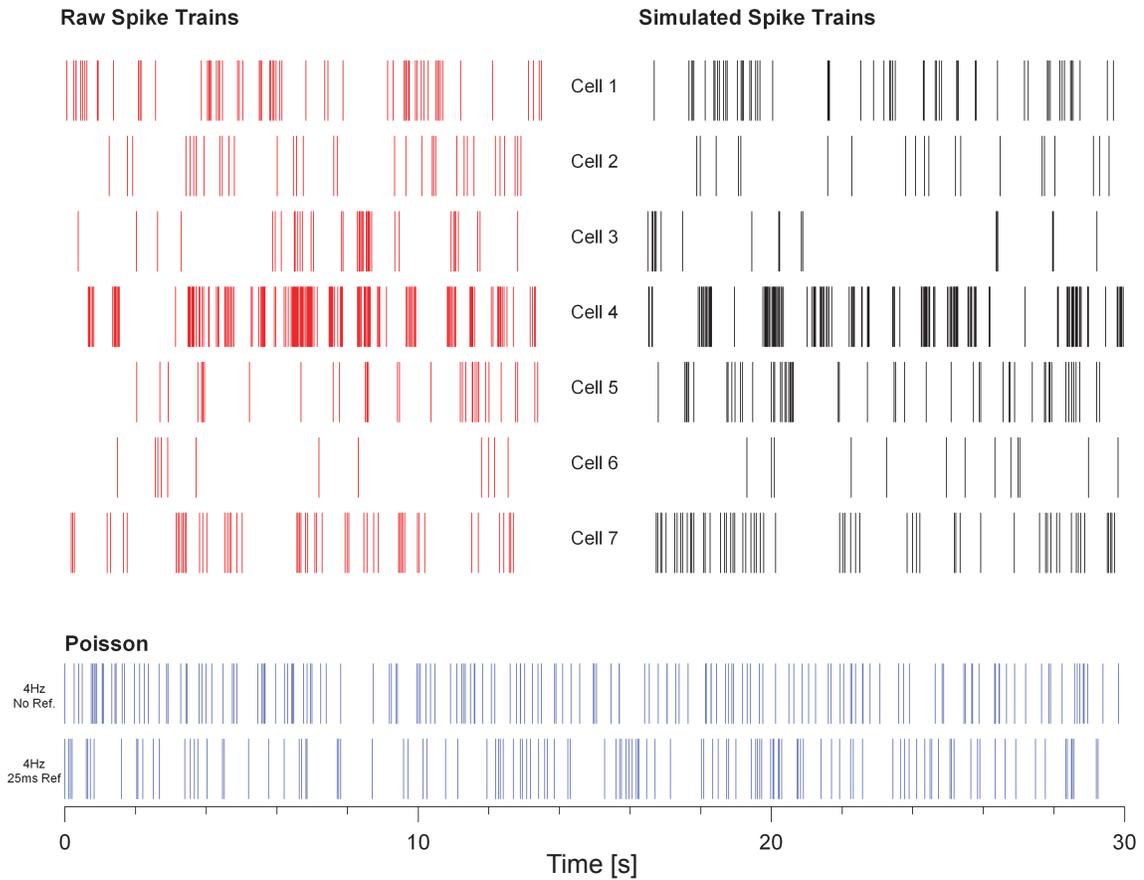


Figure A.6. Comparison of raw and simulated PN spike trains. On the left (in red) are simultaneously recorded, 12 s long examples of spontaneous spiking from the same seven PNs from figure A.5. On the right (in black) are simulated spike rasters generated from the hazard model for each of the seven PNs. Note that qualitatively similar firing patterns (including bursts) can be generated from a renewal process model (i.e., where successive ISI are completely independent). For comparison, 30 s long rasters from a Poisson model and a Poisson model with a refractory period are shown at the bottom (in blue).

PNs from figure A.4 are plotted in figure A.6. In most cases, the simple renewal model (with five parameters) is sufficient to replicate the qualitative behavior of the original PN, including its degree of burstiness.

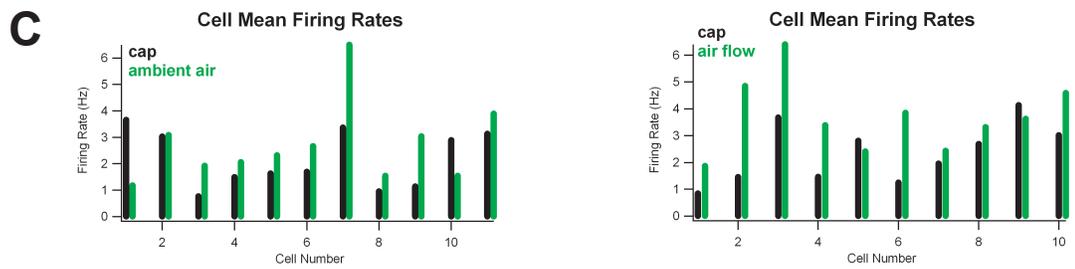
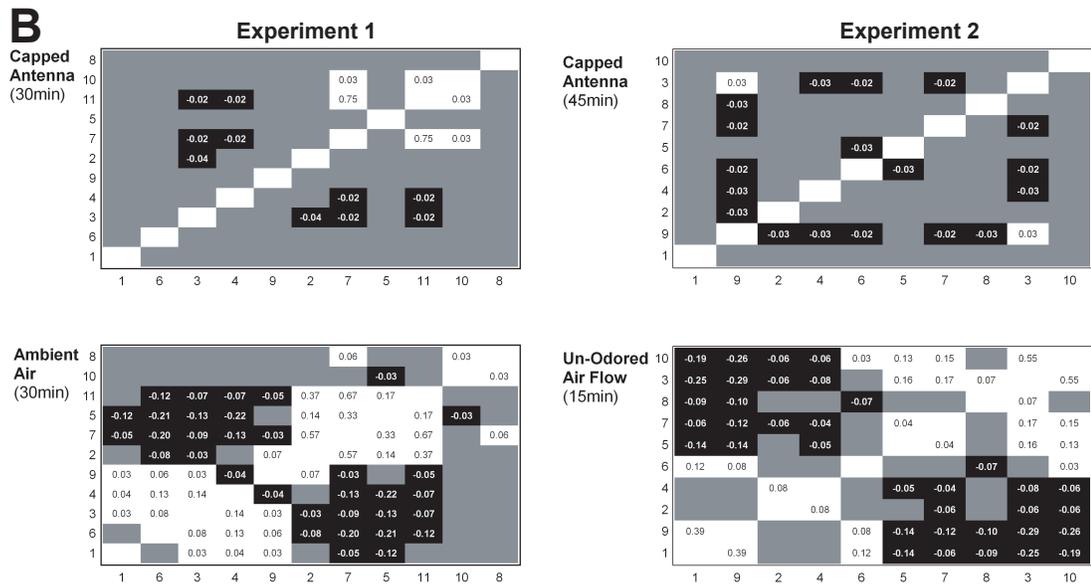
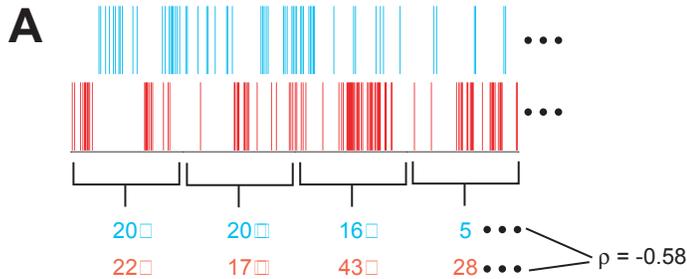
A.1.3 Measuring PN-PN correlations

The previous two sections were concerned with the firing statistics of single PNs. In this section we will examine the statistics of pairs of PNs recorded simultaneously and test whether some pairs of neurons show correlated firing at baseline.

Given two PNs recorded simultaneously, we tested their degree of correlated firing using a simple method (see figure A.7A for a diagram). We began by splitting the spike trains from the two neurons into successive, non-overlapping time bins, and then computed the spike count for each cell in every time bin. We then calculated the correlation coefficient for these two time series of spike counts. By varying the size of the time bin, we could modify the time scale over which we test for correlations. In practice, we tried bin widths between 100 ms and several seconds. All pairs of PNs included in this analysis showed stationary firing statistics (as measured by the stationarity test from section A.1.1).

To test the significance of these correlations, the time series of spike counts were randomly shuffled and the correlation coefficient recomputed. This procedure was

Figure A.7 (on the next page). Baseline correlations between simultaneously recorded PNs. *A*, schematic depicting the method of computing PN-PN correlations. Spike counts are computed for two PNs in successive, non-overlapping bins. The correlation coefficient for these two time series of spike counts is then computed. Typical bin size is between 100 milliseconds and several seconds. *B*, matrices of correlation coefficients (250 ms bins) for all pairs of cells from two different experiments. White boxes indicate a significant positive correlation, black signifies negative, and grey boxes are non-significant values (non-significant values are not shown). In each experiment, correlations were measured under two conditions: when the antenna was covered by an airtight cap; and when the antenna was free (in experiment 2, there was unodorized air flowing over the antenna). In both experiments, removing the cap resulted in a significant increase in the number of correlated pairs as well as the mean correlation strength. *C*, mean firing rate for each PN, computed separately for each condition (capped antenna versus air flow/ambient air).



repeated 100 times to generate a distribution of shuffled correlation values. The true (non-shuffled) correlation value was considered significant if it fell outside the central 95% of the shuffled distribution.

Figure A.7B shows four matrices of correlation coefficients computed across populations of PNs from two separate experiments (250 ms bins). In each matrix, the correlation coefficient for every pair of cells is computed (although only significant correlation values are shown). Cell pairs with significant positive correlations are indicated by a white box, and negative correlation by a black box. For each experiment, the correlations for two conditions are plotted: the top row shows correlations when the antenna is completely enclosed in an airtight plastic cap; in the bottom row, the cap has been removed and the antenna is exposed to room air (or a constant air flow, in experiment 2).

The data in figure A.7B indicate that in ambient air, there are often many significant positive and negative correlations between cell pairs at a time scale of 250 ms. In fact, PNs seem to form groups of cells that are all positively correlated within a group, and negatively correlated across groups.

When the antenna is sealed within an airtight cap, however, the majority of these correlations go away, and those that remain are typically weaker. This suggests that the majority of correlations observed in ambient air are due to external stimuli (presumably olfactory and/or mechanical). We did observe that in experiment 1, one pair of cells (PNs 7 and 11) showed a strong positive correlation coefficient of ~ 0.7 under both conditions.

Qualitatively similar results were found using time bin widths from 100 ms to 5 s.

A.2 Discussion

A.2.1 Modeling PN firing statistics

In this study, we began by measuring a number of statistics of individual PN spike trains, recorded in the absence of odor stimulation. We showed that we can routinely measure the ISI density (and therefore the hazard function as well) of multiple simultaneously recorded PNs. Furthermore we can restrict this analysis to periods of time when the firing statistics of all the PNs are relatively stationary (by our measure).

Next, we showed that a simple hazard model with five parameters is sufficient to explain the empirical hazard functions of all PNs recorded in this study. Although PN firing statistics are not completely consistent with a renewal point process, we were able to model them as such and generate simulated spike trains from their model hazard functions. These simulated spike trains exhibited the same qualitative degree of burstiness as the PNs true spike trains. This suggests that the majority of bursty PN behavior can be statistically explained by a single renewal process, without requiring PNs to have multiple states—each with its own firing statistics (e.g., Sherman, 2001).

STATIONARITY OF THE DATA

One set of questions that arise from this study relate to non-stationary data: what are its underlying causes, and how can we quantify and model it? The spiking activity of a large number of PNs is rarely stationary throughout a recording session of one hour or longer. In the current study, we restricted ourselves to periods of relatively stationary neural activity before computing any other statistic. The reason that the statistics of PN firing change over time is not known. Some changes may be due to external stimuli (e.g., slight changes in the air surrounding the antenna), even though we attempted to minimize this possibility. These changes, because they are stimulus evoked, are not pertinent to understanding how PNs behave at rest (e.g., with constant input).

More interesting, however, are those changes in PN firing statistics that are intrinsic to the brain. Whether such modulations do occur is not known, and neither is their underlying cause (e.g., changes in neuromodulator levels, neural plasticity). Further work will be required to determine the extent of intrinsic nonstationarities and to incorporate them into the present statistical model.

MODELING THE STATISTICS OF ODOR RESPONSES

A related problem arises when one considers how to extend the present analysis to regions of odor-evoked activity. As demonstrated in chapter 4, PN activity is far from stationary during an odor response (except for the fixed-point response to prolonged odors). The activity of PNs changes quickly—at a timescale of 150–200 ms—in response to the onset and the offset of an odor. This means that the activity of

a single PN can only be considered stationary for ~ 200 ms. Unfortunately, accurate estimates of the statistics described in this appendix require many seconds (if not minutes) of stationary data. One possibility for overcoming this problem would be to determine if any properties of baseline firing remain constant throughout the odor response for individual PNs. For example, if three or four of the parameters in our hazard model remained constant throughout an odor response, then estimating the remaining parameters would require less data. Unfortunately, determining if any parameters remain constant is a difficult task in itself.

A.2.2 PN-PN correlations

In this study, we also examined correlations in PN spiking activity under baseline conditions similar to the pre-odor periods in chapters 3 and 4 (i.e., ambient room air or a constant unodorized air flow). We found that some pairs of PNs can exhibit significant positive or negative correlations under these conditions. These correlations tended to be at a timescale of hundreds of milliseconds or longer. A separate set of tests (data not shown) were unable to find strong correlations at much shorter time scales. These correlations seem to be primarily the result of common input to the PN population, since the number and strength of significant correlations decreases substantially when the antenna is sealed in an airtight cap.

These results begin to address a number of open questions about the connectivity of PNs in the locust, and the functional strength of these connections. During an odor response, correlated PN activity is caused in part by common input from

olfactory receptor neurons (ORNs). At rest, ORNs are much less active, so correlations in PN spiking activity are likely to arise from direct or indirect connections within the antennal lobe. The degree of PN correlations that we observe suggest that functional connections between PNs are relatively weak. They seem to have minimal effect on baseline firing behavior, and any correlations that are observed occur over long timescales (hundreds of milliseconds)—too slow to strongly influence the precise timing of spikes.

A.3 Acknowledgements

Some of the work in this appendix was done in collaboration with Christophe Pouzat.

References

- ABBOTT, L. and P. DAYAN (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
- ABELES, M. (1982). Role of the cortical neuron: Integrator or coincidence detector? *Isr J Med Sci* 18:83–92.
- ABELES, M. (1991). *Corticonics: Neural Circuits of the Cerebral Cortex*. New York: Cambridge University Press.
- ABELES, M., H. BERGMAN, E. MARGALIT, and E. VAADIA (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* 70:1629–38.
- ADRIAN, E. (1926). The impulses produced by sensory nerve endings. *J Physiol (Lond)* 61:49–72.
- ADRIAN, E. (1942). Olfactory reactions in the brain of a hedgehog. *J Physiol (Lond)* 100:459–73.
- AFRAIMOVICH, V. S., V. P. ZHIGULIN, and M. I. RABINOVICH (2004). On the origin of reproducible sequential activity in neural circuits. *Chaos* 14(4): 1123–1129.
- ALLISON, A. C. (1953). The structure of the olfactory bulb and its relationship to the olfactory pathways in the rabbit and the rat. *J Comp Neurol* 98(2): 309–353.
- ATIYA, A. (1992). Recognition of multi-unit neural signals. *IEEE Transact Biomed Eng* 39:723–729.
- BARLOW, H. (1969). Pattern recognition and the responses of sensory neurones. *Ann NY Acad Sci* 156:872–81.
- BIERNACKI, C. and G. GOVAERT (1998). Choosing models in model-based clustering and discriminant analysis. *INRIA Research Report #3509:1–14*. <http://www.inria.fr/rrrt/rr-3509.html>.
- BISHOP, C. (1995). *Neural Networks for Pattern Recognition*. Oxford: Clarendon Press.
- BOYLES, R. (1983). On the convergence of the EM algorithm. *J Roy Stat Soc B* 45:47–50.
- BRAGIN, A., G. JANDÓ, Z. NÁDASDY, J. HETKE, K. WISE, and G. BUZSÁKI (1995). Gamma (40–100 Hz) oscillation in the hippocampus of the behaving rat. *J Neurosci* 15:47–69.
- BRANDT, S. (1999). *Data Analysis*. New York: Springer-Verlag.
- BRÄUNIG, P. (1991). Suboesophageal DUM neurons innervate the principal neuropiles of the locust brain. *Phil Trans R Soc Lond B* 332:221–240.
- BRECHT, M., M. SCHNEIDER, B. SAKMANN, and T. W. MARGRIE (2004). Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature* 427(6976): 704–710.

- BUZSAKI, G. and A. DRAGUHN (2004). Neuronal oscillations in cortical networks. *Science* 304(5679): 1926–1929.
- CAIN, W. and B. POTTS (1996). Switch and bait: probing the discriminative basis of odor identification via recognition memory. *Chem Senses* 21:35–44.
- CARR, C. E. and M. KONISHI (1990). A circuit for detection of interaural time differences in the brain stem of the barn owl. *J Neurosci* 10(10): 3227–3246.
- CHAPMAN, R. (1998). *The Insects: Structure and Function* (4th ed.). Cambridge: Cambridge University Press.
- CLYNE, P. J., S. J. CERTEL, M. DE BRUYNE, L. ZASLAVSKY, W. A. JOHNSON, and J. R. CARLSON (1999). The odor specificities of a subset of olfactory receptor neurons are governed by *acj6*, a pou-domain transcription factor. *Neuron* 22(2): 339–47.
- CONTRERAS, D., A. DESTEXHE, and M. STERIADE (1997). Intracellular and computational characterization of the intracortical inhibitory control of synchronized thalamic inputs in vivo. *J Neurophysiol* 78:335–30.
- CSIBRA, G., G. DAVIS, M. W. SPRATLING, and M. H. JOHNSON (2000). Gamma oscillations and object processing in the infant brain. *Science* 290(5496): 1582–5.
- DAYAN, P. and L. F. ABBOTT (2001). *Theoretical Neuroscience*. Cambridge, MA: MIT Press.
- DE BELLE, J. S. and M. HEISENBERG (1994). Associative odor learning in *Drosophila* abolished by chemical ablation of mushroom bodies. *Science* 263(5147): 692–5.
- DE BRUYNE, M., K. FOSTER, and J. R. CARLSON (2001). Odor coding in the *Drosophila* antenna. *Neuron* 30(2): 537–552.
- DEMPSTER, A., N. LAIRD, and D. RUBIN (1977). Maximum likelihood from incomplete data via the EM algorithm. *J Royal Stat Soc B* 39:1–38.
- DIESMANN, M., M. O. GEWALTIG, and A. AERTSEN (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402(6761): 529–33.
- DOBRI TSA, A. A., W. VAN DER GOES VAN NATERS, C. G. WARR, R. A. STEINBRECHT, and J. R. CARLSON (2003). Integrating the molecular and cellular basis of odor coding in the *Drosophila* antenna. *Neuron* 37(5): 827–841.
- DRAKE, K., K. WISE, J. FARRAYE, D. ANDERSON, and S. BEMENT (1988). Performance of planar multi-site microprobes in recording extracellular single-unit intracortical activity. *IEEE Trans Biomed Eng* 35:719–732.
- DUBNAU, J., L. GRADY, T. KITAMOTO, and T. TULLY (2001). Disruption of neurotransmission in *Drosophila* mushroom body blocks retrieval but not acquisition of memory. *Nature* 411(6836): 476–80.
- DUCHAMP-VIRET, P., B. PALOUZIER-PAULIGNAN, and A. DUCHAMP (1996). Odor coding properties of frog olfactory cortical neurons. *Neuroscience* 74(3): 885–95.
- ECKHORN, R., R. BAUER, W. JORDAN, M. BROSCHE, W. KRUSE, M. MUNK, and H. J. REITBOECK (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol Cybern* 60(2): 121–130.

- ENGEL, A. K., P. FRIES, and W. SINGER (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2(10): 704–16.
- ERBER, J., U. HOMBERG, and W. GRONENBERG (1987). Functional roles of the mushroom bodies in insects. In *Arthropod Brain: Its Evolution, Development, Structure and Functions*, ed. A. Gupta. Wiley.
- ERNST, K.-D., J. BOECKH, and V. BOECKH (1977). A neuroanatomical study on the organization of the central antennal pathways in insects. *Cell Tiss. Res.* 176:285–308.
- FARIVAR, S. S. (2005). *Cytoarchitecture of the Locust Olfactory System*. Ph.D. thesis, California Institute of Technology, Pasadena.
- FEE, M., P. MITRA, and D. KLEINFELD (1996a). Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J Neurosci Methods* 69:175–188.
- FEE, M., P. MITRA, and D. KLEINFELD (1996b). Variability of extracellular spike waveforms of cortical neurons. *J Neurophysiol* 76:3823–3833.
- FITZPATRICK, D. C., R. BATRA, T. R. STANFORD, and S. KUWADA (1997). A neuronal population code for sound localization. *Nature* 388(6645): 871–874.
- FLANAGAN, D. and A. MERCER (1989). Morphology and response characteristics of neurones in the deutocerebrum of the brain in the honeybee *Apis mellifera*. *J Comp Physiol A* 164:483–449.
- FONTA, C., X. SUN, and C. MASSON (1993). Morphology and spatial distribution of bee antennal lobe interneurons responsive to odours. *Chem. Senses* 18:101–19.
- FRALEY, C. and A. RAFTERY (1998). How many clusters? Which clustering method? Answers via model-based cluster analysis. *Comput J* 41:578–588.
- FRICKER, D. and R. MILES (2000). EPSP amplification and the precision of spike timing in hippocampal neurons. *Neuron* 28:559–69.
- FRIEDRICH, R. W. and G. LAURENT (2001). Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. *Science* 291(5505): 889–894.
- FRIES, P., J. H. REYNOLDS, A. E. RORIE, and R. DESIMONE (2000). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–63.
- FROSTIG, R., E. LIEKE, D. TSO, and A. GRINVALD (1990). Cortical functional architecture and local coupling between neuronal activity and the microcirculation revealed by in vivo high resolution optical imaging of intrinsic signals. *PNAS* 87:6082–6086.
- GALÁN, R. F., S. SACHSE, C. G. GALIZIA, and A. V. M. HERZ (2004). Odor-driven attractor dynamics in the antennal lobe allow for simple and rapid olfactory pattern classification. *Neural Comput* 16(5): 999–1012.
- GALARRETA, M. and S. HESTRIN (2001). Spike transmission and synchrony detection in networks of GABAergic interneurons. *Science* 292(5525): 2295–9.
- GAO, Q., B. YUAN, and A. CHESS (2000). Convergent projections of *Drosophila* olfactory neurons to specific glomeruli in the antennal lobe. *Nat Neurosci* 3:780–85.
- GELPERIN, A. (1999). Oscillatory dynamics and information processing in olfactory systems. *J Exp Biol* 202:1855–1864.

- GELPERIN, A. and D. W. TANK (1990). Odour-modulated collective network oscillations of olfactory interneurons in a terrestrial mollusc. *Nature* 345:437–40.
- GIURFA, M., S. ZHANG, A. JENETT, R. MENZEL, and M. V. SRINIVASAN (2001). The concepts of ‘sameness’ and ‘difference’ in an insect. *Nature* 410(6831): 930–933.
- GLASER, E. and W. MARKS (1968). Online separation of interleaved neuronal pulse sequences. *Data Acquisition Process Biol Med* 5:137–156.
- GOLDMAN, A. L., W. VAN DER GOES VAN NATERS, D. LESSING, C. G. WARR, and J. R. CARLSON (2005). Coexpression of two functional odor receptors in one neuron. *Neuron* 45(5): 661–666.
- GRAY, C. (1994). Synchronous oscillations in neuronal systems: mechanisms and function. *J Comput Neurosci* 1:11–38.
- GRAY, C. M., P. KONIG, A. K. ENGEL, and W. SINGER (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* 338(6213): 334–337.
- GRONENBERG, W. (1999). Modality-specific segregation of input to ant mushroom bodies. *Brain Behav Evol* 54(2): 85–95.
- HABERLY, L. (1990). Comparative aspects of olfactory cortex. In *Cerebral Cortex*, eds. E. Jones and A. Peters, pp. 137–66. New York: Plenum Press.
- HAMMER, M. (1993). An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature* 366:59–63.
- HANSSON, B., S. A. OCHIENG, X. GROSMAITRE, S. ANTON, and P. G. N. NJAGI (1996). Physiological responses and central nervous projections of antennal olfactory receptor neurons in the adult desert locust, *Schistocerca gregaria* (orthoptera: Acrididae). *J. Comp. Physiol. A* 179:157–167.
- HANSSON, B. S. and S. ANTON (2000). Function and morphology of the antennal lobe: new developments. *Annu Rev Entomol* 45:203–31.
- HARRIS, K., D. HENZE, J. CSICSVARI, H. HIRASE, and G. BUZSAKI (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol* 84:401–414.
- HEISENBERG, M. (1989). Genetic approach to learning and memory (mnemogenetics) in *Drosophila melanogaster*. In *Fundamentals of Memory Formation: Neuronal Plasticity and Brain Function*, ed. B. Rahmann, pp. 3–45. New York: Gustav Fischer Verlag.
- HEISENBERG, M. (2003). Mushroom body memoir: from maps to models. *Nat Rev Neurosci* 4(4): 266–275.
- HEISENBERG, M., A. BORST, S. WAGNER, and D. BYERS (1985). *Drosophila* mushroom body mutants are deficient in olfactory learning. *J Neurogenet* 2:1–30.
- HILDEBRAND, J. G. and G. SHEPHERD (1997). Mechanisms of olfactory discrimination: converging evidence for common principles across phyla. *Annu Rev Neurosci* 20:595–631.
- HOMBERG, U., T. A. CHRISTENSEN, and J. G. HILDEBRAND (1989). Structure and function of the deutocerebrum in insects. *Annu Rev Entomol* 34:477–501.
- HOPFIELD, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–6.

- HUBEL, D. and T. WIESEL (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154.
- HULATA, E., R. SEGEV, Y. SHAPIRA, M. BENVENISTE, and E. BEN-JACOB (2000). Detection and sorting of neural spikes using wavelet packets. *Phys Rev Letters* 85:4637–4640.
- IKEGAYA, Y., G. AARON, R. COSSART, D. ARONOV, I. LAMPL, D. FERSTER, and R. YUSTE (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304(5670): 559–564.
- JACK, J., S. REDMAN, and K. WONG (1981). The components of synaptic potentials evoked in cat spinal motoneurons by impulses in single group Ia afferents. *J Physiol (Lond)* 321:65–96.
- JOHANSSON, R. S. and I. BIRZNIKS (2004). First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nat Neurosci* 7(2): 170–177.
- JOHNSON, D. H. (1996). Point process models of single-neuron discharges. *J Comput Neurosci* 3(4): 275–299.
- JORTNER, R. A., S. S. FARIVAR, and G. LAURENT. Dense connectivity for sparse representation in an olfactory network. Submitted.
- KAISLING, K.-E. (1971). Insect olfaction. In *Handbook of Sensory Physiology*, ed. L. Beidler, Volume 4, pp. 351–431. Berlin: Springer.
- KAISLING, K.-E. (1995). Single unit and electroantennogram recordings in insect olfactory organs. In *Experimental Cell Biology of Taste and Olfaction*, eds. A. I. Spielman and J. G. Brand. New York: CRC Press.
- KANERVA, P. (1988). *Sparse Distributed Memory*. Cambridge: MIT Press.
- KENYON, F. (1896). The brain of the bee. *J Comp Neurol* 6:133–210.
- KÖNIG, P., A. K. ENGEL, and W. SINGER (1996). Integrator of coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci* 19:130–37.
- KREITER, A. K. and W. SINGER (1996). Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J Neurosci* 16(7): 2381–2396.
- LAISSUE, P. P., C. REITER, P. R. HIESINGER, S. HALTER, K. F. FISCHBACH, and R. F. STOCKER (1999). Three-dimensional reconstruction of the antennal lobe in *Drosophila melanogaster*. *J Comp Neurol* 405(4): 543–552.
- LAM, Y. W., L. B. COHEN, M. WACHOWIAK, and M. R. ZOCHOWSKI (2000). Odors elicit three different oscillations in the turtle olfactory bulb. *J Neurosci* 20(2): 749–62.
- LARSSON, M. C., A. I. DOMINGOS, W. D. JONES, M. E. CHIAPPE, H. AMREIN, and L. B. VOSSHALL (2004). Or83b encodes a broadly expressed odorant receptor essential for *Drosophila* olfaction. *Neuron* 43(5): 703–714.
- LAURENT, G. (1999). A systems perspective on early olfactory coding. *Science* 286(5440): 723–8.
- LAURENT, G. and H. DAVIDOWITZ (1994). Encoding of olfactory information with oscillating neural assemblies. *Science* 265:1872–1875.
- LAURENT, G. and M. NARAGHI (1994). Odorant-induced oscillations in the mushroom bodies of the locust. *J Neurosci* 14:2993–3004.

- LAURENT, G., M. STOPFER, R. W. FRIEDRICH, M. I. RABINOVICH, A. VOLKOVSKII, and H. D. ABARBANEL (2001). Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu Rev Neurosci* 24:263–97.
- LAURENT, G., M. WEHR, and H. DAVIDOWITZ (1996). Temporal representations of odors in an olfactory network. *J Neurosci* 16(12): 3837–47.
- LEITCH, B. and G. LAURENT (1996). GABAergic synapses in the antennal lobe and mushroom body of the locust olfactory system. *J Comp Neurol* 372(4): 487–514.
- LETELIER, J. and P. WEBER (2000). Spike sorting based on discrete wavelet transform coefficients. *J Neurosci Methods* 101:93–106.
- LEWICKI, M. (1994). Bayesian modeling and classification of neural signals. *Neural Computation* 6:1005–1030.
- LEWICKI, M. (1998). A review of methods for spike sorting: the detection and classification of neuronal action potentials. *Network* 9:R53–R78.
- LI, Y. and N. J. STRAUSFELD (1999). Multimodal efferent and recurrent neurons in the medial lobes of cockroach mushroom bodies. *J Comp Neurol* 409(4): 647–663.
- LING, L. and D. TOLHURST (1983). Recovering the parameters of finite mixtures of normal distributions from noisy record: an empirical comparison of different estimating procedures. *J Neurosci Meth* 8:309–333.
- LINSTER, C. and B. SMITH (1999). Generalization between binary odor mixtures and their components in the rat. *Physiol Behav* 66:701–7.
- LIVERMORE, A. and D. LAING (1996). Influence of training and experience on the perception of multicomponent odor mixtures. *J Exp Psychol* 22:267–77.
- MACLEOD, K., A. BACKER, and G. LAURENT (1998). Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature* 395(6703): 693–8.
- MACLEOD, K. and G. LAURENT (1996). Distinct mechanisms for synchronization and temporal patterning of odor-encoding neural assemblies. *Science* 274(5289): 976–9.
- MALDONADO, P. E., S. FRIEDMAN-HILL, and C. M. GRAY (2000). Dynamics of striate cortical activity in the alert macaque: II. Fast time scale synchronization. *Cereb Cortex* 10(11): 1117–1131.
- MALUN, D., U. WALDOW, D. KRAUS, and J. BOECKH (1993). Connections between the deutocerebrum and the protocerebrum, and neuroanatomy of several classes of deutocerebral projection neurons in the brain of male *Periplaneta americana*. *J Comp Neurol* 329(2): 143–162.
- MCGUIRE, S. E., P. T. LE, and R. L. DAVIS (2001). The role of *Drosophila* mushroom body signaling in olfactory memory. *Science* 293(5533): 1330–3.
- McLACHLAN, G. and T. KRISHNAN (1997). *The EM Algorithm and Extensions*. New York: John Wiley & Sons.
- MENZEL, R. and U. MULLER (1996). Learning and memory in honeybees: from behavior to neural substrates. *Annu Rev Neurosci* 19:379–404.
- MILLECCHIA, R. and R. McINTYRE (1978). Automatic nerve impulse identification and separation. *Comput Biomed Res* 11:459–468.

- MOMBAERTS, P., F. WANG, C. DULAC, S. K. CHAO, A. NEMES, M. MENDELSON, J. EDMONDSON, and R. AXEL (1996). Visualizing an olfactory sensory map. *Cell* 87(4): 675–86.
- OCHIENG, S. A., E. HALLBERG, and B. S. HANSSON (1998). Fine structure and distribution of antennal sensilla of the desert locust, *Schistocerca gregaria* (Orthoptera: Acrididae). *Cell Tissue Res* 291(3): 525–536.
- OGAWA, S., D. TANK, R. MENON, J. ELLERMANN, S. KIM, H. MERKLE, and K. UGURBIL (1992). Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *PNAS* 89:5951–5955.
- O'KEEFE, J. and M. L. RECCE (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3(3): 317–330.
- OLSKEVICH, S. (1999). Cholinergic synaptic transmission in insect mushroom bodies in vitro. *J Neurophysiol* 82(2): 1091–1096.
- PAPOULIS, A. (1980). *Circuits and Systems. A Modern Approach*. New York: Holt, Rinehart and Winston.
- PARKER, A. J. and W. T. NEWSOME (1998). Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227–277.
- PATEL, A. D. and E. BALABAN (2000). Temporal patterns of human cortical activity reflect tone sequence structure. *Nature* 404:80–84.
- PEREZ-ORIVE, J. (2004). *Neural oscillations and the decoding of sensory information*. Ph.D. thesis, California Institute of Technology, Pasadena.
- PEREZ-ORIVE, J., M. BAZHENOV, and G. LAURENT (2004). Intrinsic and circuit properties favor coincidence detection for decoding oscillatory input. *J Neurosci* 24(26): 6037–47.
- PEREZ-ORIVE, J., O. MAZOR, G. C. TURNER, S. CASSENAER, R. I. WILSON, and G. LAURENT (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science* 297(5580): 359–65.
- PETERSEN, R. S., S. PANZERI, and M. E. DIAMOND (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron* 32(3): 503–514.
- POUGET, A., S. DENEVE, J. C. DUCOM, and P. E. LATHAM (1999). Narrow versus wide tuning curves: What's best for a population code? *Neural Comput* 11(1): 85–90.
- POUILLE, F. and M. SCANZIANI (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293(5532): 1159–63.
- POUZAT, C., M. DELESCLUSE, P. VIOT, and J. DIEBOLT (2004). Improved spike-sorting by modeling firing statistics and burst-dependent spike amplitude attenuation: a Markov chain Monte Carlo approach. *J Neurophysiol* 91(6): 2910–2928.
- POUZAT, C., O. MAZOR, and G. LAURENT (2002). Using noise signature to optimize spike-sorting and to assess neuronal classification quality. *J Neurosci Methods* 122(1): 43–57.
- RABINOVICH, M., A. VOLKOVSKII, P. LECANDA, R. HUERTA, H. D. ABARBANEL, and G. LAURENT (2001). Dynamical encoding by networks of competing neuron groups: winnerless competition. *Phys Rev Lett* 87(6): 068102.
- REDNER, R. and H. WALKER (1984). Mixture densities, maximum likelihood and the EM algorithm. *SIAM Rev* 26:195–239.

- RIEHLE, A., S. GRÜN, M. DIEMANN, and A. AERTSEN (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 281:34–42.
- RIEKE, F., D. WARLAND, R. DE RUYTER VAN STEVENINCK, and W. BIALEK (1997). *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press.
- RINGACH, D. L., M. J. HAWKEN, and R. SHAPLEY (1997). Dynamics of orientation tuning in macaque primary visual cortex. *Nature* 387(6630): 281–284.
- RODRIGUEZ, R., M. GEORGE, J. P. LACHAUX, J. MARTINERIE, B. RENAULT, and F. J. VARELA (1999). Perception's shadow: long-distance synchronization of human brain activity. *Nature* 397:430–433.
- ROEDER, K. D. (1948). Organization of the ascending giant fiber system in the cockroach (*Periplaneta americana* L.). *J Exper Zool* 108:327–338.
- ROLLS, E. T. and M. J. TOVEE (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. *J Neurophysiol* 73(2): 713–726.
- SAHANI, M. (1999). *Latent variable models for neural data analysis*. Ph.D. thesis, California Institute of Technology, Pasadena.
- SCHAFER, S., J. ROSENBLUM, and R. MENZEL (1994). Ionic currents of kenyon cells from the mushroom body of the honeybee. *J Neurosci* 14:4600–12.
- SCHNITZER, M. J. and M. MEISTER (2003). Multineuronal firing patterns in the signal from eye to brain. *Neuron* 37(3): 499–511.
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann Stat* 6:461–464.
- SHADLEN, M. N. and J. A. MOVSHON (1999). Synchrony unbound: A critical evaluation of the temporal binding hypothesis. *Neuron* 24:67–77.
- SHEPHERD, G., W. CHEN, and C. GREER (2004). Olfactory bulb. In *The Synaptic Organization of the Brain*, ed. G. Shepherd. New York: Oxford University Press.
- SHEPHERD, G. M. (1972). Synaptic organization of the mammalian olfactory bulb. *Physiol Rev* 52(4): 864–917.
- SHERMAN, S. M. (2001). Tonic and burst firing: dual modes of thalamocortical relay. *Trends Neurosci* 24(2): 122–126.
- SIAPAS, A. G., E. V. LUBENOV, and M. A. WILSON (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron* 46(1): 141–151.
- SINGER, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24(1): 49–65.
- SINGER, W. and C. GRAY (1995). Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–86.
- SKAGGS, W. E., B. L. McNAUGHTON, M. A. WILSON, and C. A. BARNES (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6(2): 149–172.
- SPITZER, H., R. DESIMONE, and J. MORAN (1988). Increased attention enhances both behavioral and neuronal performance. *Science* 240(4850): 338–340.

- STERIADE, M., F. AMZICA, and D. CONTRERAS (1996). Synchronization of the fast (30–40 Hz) spontaneous cortical rhythms during brain activation. *J Neurosci* 16:392–417.
- STOPFER, M., S. BHAGAVAN, B. H. SMITH, and G. LAURENT (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature* 390(6655): 70–4.
- STOPFER, M., V. JAYARAMAN, and G. LAURENT (2003). Intensity versus identity coding in an olfactory system. *Neuron* 39(6): 991–1004.
- STOPFER, M. and G. LAURENT (1999). Short-term memory in olfactory network dynamics. *Nature* 402(6762): 664–8.
- STRAUSFELD, N. J., L. HANSEN, Y. LI, R. S. GOMEZ, and K. ITO (1998). Evolution, discovery, and interpretations of arthropod mushroom bodies. *Learn Mem* 5(1-2): 11–37.
- THEUNISSEN, F. and J. P. MILLER (1995). Temporal encoding in nervous systems: a rigorous definition. *J Comput Neurosci* 2(2): 149–162.
- THORPE, S. J. (1990). Spike arrival times: A highly efficient coding scheme for neural networks. In *Parallel Processing in Neural Systems*, eds. R. Eckmiller, G. Hartmann, and G. Hauske, pp. 91–94. Amsterdam: Elsevier.
- THORPE, S. J. (2001). Spike-based strategies for rapid processing. *Neural Netw* 14:715–725.
- VANRULLEN, R., R. GUYONNEAU, and S. J. THORPE (2005). Spike times make sense. *Trends Neurosci* 28(1): 1–4.
- VINJE, W. E. and J. L. GALLANT (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287(5456): 1273–1276.
- VON DER MALSBURG, C. and W. SCHNEIDER (1986). A neural cocktail-party processor. *Biol Cybern* 54(1): 29–40.
- VOSSHALL, L. B., H. AMREIN, P. S. MOROZOV, A. RZHETSKY, and R. AXEL (1999). A spatial map of olfactory receptor expression in the *Drosophila* antenna. *Cell* 96(5): 725–36.
- VOSSHALL, L. B., A. M. WONG, and R. AXEL (2000). An olfactory sensory map in the fly brain. *Cell* 102(2): 147–59.
- WEHNER, R. and R. MENZEL (1990). Do insects have cognitive maps? *Annu Rev Neurosci* 13:403–414.
- WEHR, M. and G. LAURENT (1996). Odor encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384:162–6.
- WEHR, M., J. PEZARIS, and M. SAHANI (1999). Simultaneous paired intracellular and tetrode recordings for evaluating the performance of spike sorting algorithms. *Neurocomputing* 26-27:1061–1068.
- WICKLEIN, M. and N. STRAUSFELD (2000). Organization and significance of neurons that detect change of visual depth in the hawk moth *manduca sexta*. *J Comp Neurol* 424:356–76.
- WILLMORE, B. and D. TOLHURST (2001). Characterizing the sparseness of neural codes. *Network: Comput Neural Syst* 12:255–70.
- WONG, A. M., J. W. WANG, and R. AXEL (2002). Spatial representation of the glomerular map in the *Drosophila* protocerebrum. *Cell* 109(2): 229–241.
- WU, C. (1983). On the convergence properties of the EM algorithm. *Annals Stat* 11:95–103.

- WU, J., L. COHEN, and C. FALK (1994). Neuronal activity during different behaviors in *Aplysia*: a distributed organization? *Science* 263:820–823.
- ZARS, T., M. FISCHER, R. SCHULZ, and M. HEISENBERG (2000). Localization of a short-term memory in *Drosophila*. *Science* 288(5466): 672–5.
- ZHU, Z., K. LIN, and T. KASAMATSU (2002). Artifactual synchrony via capacitance coupling in multi-electrode recording from cat striate cortex. *J Neuroscience Methods* 115:45–53.
- ZOU, Z., L. F. HOROWITZ, J. P. MONTMAYEUR, S. SNAPPER, and L. B. BUCK (2001). Genetic tracing reveals a stereotyped sensory map in the olfactory cortex. *Nature* 414(6860): 173–9.