Sensitivity Enhancement Mechanisms at the Periphery of the Olfactory Pathway

Thesis submitted for the degree of

Doctor of Philosophy

at the University of Leicester

by

Josephine Ann Mackenzie

Department of Engineering

University of Leicester

January 2007

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Massive convergence of input from olfactory receptor neurons (ORNs) with identical tunings leads to spatial integration of sensory signals, thereby boosting sensitivity to sensory cues. The consequent reduction in detection thresholds is assumed to derive from the pooling of elevated firing rates across the ORN population. By comparing detection thresholds at the first two stages of the olfactory pathway in an olfactory specialist, the moth, allowed for the quantification of the sensitivity boost achieved during early sensory processing. This boost was found to be at least 3 orders of magnitude, which was shown to exceed that achieved by a theoretical model of spike train integration. The sensitivity enhancement achieved by this system therefore goes beyond straightforward spatial summation of receptor firing rates, suggesting subtler coding and readout mechanisms.

To discount the possibility of ORNs employing a temporal encoding scheme, an investigation into spike patterns at the periphery was performed. While no temporal patterns were evident, a temporal encoding scheme remains a possibility. Despite the inconclusive result found here, the analysis demonstrates the need for an investigation of the stationarity of spike trains, where a statistical basis underlies the analysis method, before drawing conclusions.

Regardless of the encoding scheme employed at the periphery, due to the multitude of possible synaptic connections within a glomerulus, it seems unlikely that this site of convergence of receptor input would be passive. A simple, but biologically plausible computational model was developed, where specific zones of the dendritic tree of an output neuron form

individual subunits capable of performing a nonlinear threshold function on ORN inputs. This nonlinear model consistently outperformed a comparable linear model when assessing the stimulus detection performance of the output neuron.

Acknowledgements

There are various people who have helped with the creation of this thesis and whom I'd like to thank:

JT for his constant support, for being an invaluable sounding board, and for his abilities in the kitchen.

My family for their motivation, encouragement and support.

The members of the Engineering Department, in particular, all of the 5th floor lab occupants, for their friendship and constructive conversations.

The European Commission for the funding of this project, and to collaborators at the Sveriges Lantbruksuniversitet in Sweden for their electrophysiological recordings, without which this work could not have been completed.

And of course, Tim Pearce for his supervision and inspiration throughout.

Contents

List of Abbreviations

AL	Antennal Lobe
ALN	Antennal Lobe Neuron
AUC	Area Under (ROC) Curve
CV	Coefficient of Variation
EAG	Electroantennogram
FFT	Fast Fourier Transform
ISF	Instantaneous Spike Frequency
ISIs	InterSpike Intervals
jPSTH	joint Peri-Stimulus Time Histogram
LNs	Local Interneurons
LSM	Linear Synapse Model
MGC	Macroglomerular Complex
NSM	Nonlinear Subunit Model
OB	Olfactory Bulb
OBPs	Odorant Binding Proteins
ORNs	Olfactory Receptor Neurons
PC	Principal Component
PCA	Principal Components Analysis
PNs	Projection Neurons
PSTHs	Peri-Stimulus Time Histograms
ROC	Receiver Operating Characteristic Curve
SE	Standard Error
Syntech	Syntech, Hilversum, The Netherlands
TES	*N*-Tris-methyl-2-aminoethanesulfonic acid
UEA	Unitary Events Analysis
UEMWA	Unitary Events by Moving Window Analysis
WPGMA	Weighted Pair-Group Method using Arithmetic Averages
Z9, E11-14:OAc	(Z, E)-9, 11-tetradecadienyl acetate

List of Publications

While no papers have yet been published in relation to this work, there have been the following poster presentations, and manuscripts currently in preparation.

1. Mackenzie, J.A., Han, Q., Takanashi, T., Skals, N., Pearce, T.C. and Hansson, B.S. (2004). Sensitivity enhancement in early olfactory processing in the moth *Spodoptera littoralis*. 7th International Congress of Neuroethology, Nyborg, Denmark. (poster).

2. Mackenzie, J.A., Han, Q., Takanashi, T., Skals, N., Bäckman, A.-C., Zimerli, L., Verschure, P.F.M.J., Pearce, T.C. and Hansson, B.S. (2005). Olfactory sensitivity enhancement beyond pooling of firing rates across sensory neurons. Computational and Systems Neuroscience, Salt Lake City, Utah. (poster).

3. Mackenzie, J.A., Han, Q., Takanashi, T., Skals, N., Bäckman, A.-C., Zimerli, L., Verschure, P.F.M.J., Pearce, T.C. and Hansson, B.S. (2007). Olfactory sensitivity enhancement beyond pooling of firing rates across sensory neurons. (manuscript in preparation).

4. Mackenzie, J.A., Takanashi, T., Skals, N., Pearce, T.C. and Hansson, B.S. (2007). Preliminary investigation into receptor neuron firing reveals differing frequency dynamics. (manuscript in preparation).

Chapter 1

Introduction

1.1 Aim of the Thesis

The initial aim of this thesis was to investigate the stimulus detection thresholds at the first two stages of the olfactory pathway of a moth. As time progressed and results indicated a signal boost between these stages that was higher than expected, the aim shifted more toward investigating how this boost could be achieved. The main emphasis then transferred to the development of a computational model used to explore several possible hypotheses.

1.2 Overview of Moth Olfaction

For many animals, olfactory cues play an important role in various activities, in particular, reproduction and the search for food. Moths have been shown to be a useful model for the investigation of olfaction, both behaviourally and using electrophysiology (for reviews see Hansson, 1995, 2002; Mustaparta, 1996). Female moths are known to use plant-produced odours to determine the suitability of sites for feeding and oviposition (Anderson et al., 1993; Anton and Hansson, 1994; Jönsson and Anderson, 1999) while insect-released odours are used for con-specific communication, often by means of sexual pheromones (Anderson et al., 1993; Anton and Hansson, 1994, 1995; Mustaparta, 1996). A stationary female usually releases these sex pheromones into the air, and it is the male that must detect and find the female by following the odour plume, often over large distances (up to 4000m away

for some moths, with favourable environmental and wind conditions (Wilson, 1963)). In order to achieve this task, the pheromone detection system of the male moth has developed into a highly specialised subsystem of the olfactory pathway (Anton and Hansson, 1995; Kalinovà et al., 2001; Mankin et al., 1991; Mustaparta, 1996). It is so finely tuned, that the olfactory system of *Bombyx Mori* is thought to require only a single pheromone molecule to trigger a nerve impulse in the male (Kaissling, 1971; Kaissling and Priesner, 1970; Minor and Kaissling, 2003), and as few as six pheromone molecules reaching the antenna are needed to alter the heart beat rhythm of the male *Spodoptera littoralis* (Angioy et al., 2003). The remainder of this chapter provides a brief overview of the moth olfactory system, from odour detection at the antenna, to the first levels of processing in the antennal lobe.

1.2.1 Structure of the First Stages of the Moth Olfactory Pathway

1.2.1.1 Odour Reception at the Sensilla

The very first stage of moth olfaction occurs in the sensilla, the thousands of tiny hair-like structures covering the surface of the antennae. It is these sensilla that house the neurons capable of detecting odour molecules. Each sensillum generally consists of one or more bipolar ORNs at its base, surrounded by supporting cells, the details of which are described by Keil (1999). Each ORN consists of a primary (or ciliary) dendrite extending into the sensillum, and a long axon, which extends along the antennal nerve and into the AL (see Figure 1.1 a)). The surface of each sensillum contains many pores through which the odour molecules diffuse. These molecules attach to odorant binding proteins (OBPs) that have been hypothesised to transport them through the aqueous sensillum lymph until they reach a receptor molecule on the ciliary dendrite (Lerner et al., 1990; Van den Berg and Ziegelberger, 1991; see Figure 1.1 b)). Activation of the highly specific receptors causes ion channels to open, leading to the generation of a dendritic current, which travels along the dendrite to the soma of the ORN. This activation of the receptor site may also cause the OBP to change so the odorant can no longer activate the receptor (Kaissling, 1986b;

Kasang, 1973). Enzymes in the sensillum lymph then break down the odour molecules into inactive metabolites (see Stengl et al., 1999 for a detailed description of the peri-reception process). If the electrical signal received by the soma of an ORN is sufficient to raise the membrane potential above a threshold voltage, an action potential (or spike) is generated, which then travels down the axon to the AL.

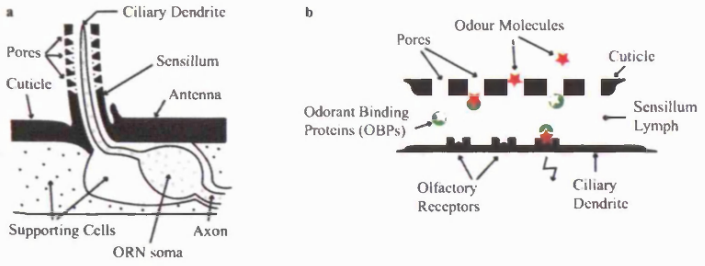

Figure 1.1: Odour reception at the sensilla. a) Close up view of a single sensillum containing a single ORN at the base with a single ciliary dendrite extending into the sensillum. The surface of the sensillum contains many pores through which odour molecules can pass. Supporting cells surround the ORN, but these are not shown in detail. (After Keil and Steinbrecht, 1987). b) Schematic diagram of a sensillum. Odour molecules in the air pass through the pores into the aqueous sensillum lymph. Here, OBPs attach to the molecules and transport them to the ciliary dendrite. The odour molecules bind with receptors on the surface of the dendrite. This binding causes an electrical signal to travel along the dendrite and into the soma of the ORN. In some cases this binding causes a change in the OBP so it can temporarily no longer transport odour molecules to the dendrite. (After Stengl et al., 1999).

1.2.2 Structure of the Antennal Lobe

From the antenna, the ORN axons travel down the antennal nerve into the ipsilateral AL (Hansson et al., 1992; Koontz and Schneider, 1987). The AL consists largely of spheroids of tightly packed neuropil (dendritic and axonal branches) known as glomeruli. Moths have around 60 of these glomeruli (Rospars, 1983; Rospars and Hildebrand, 1992), each of which are innervated by ORNs, local interneurons (LNs) and projection neurons (PNs). In general, LNs are able to innervate any number of glomeruli, and PNs act as the output neurons of the AL. ORNs are thought to house only a single receptor type, and axons from functionally characterised ORNs have been found to project in a chemotypic manner to single identified glomeruli (Berg et al., 1998; Hansson et al., 1995, 1992; Ochieng' et al., 1995; Todd et al.,

1995; see Figures 1.2 and 1.3). Synaptic connections within the AL are almost entirely restricted to within the glomeruli.

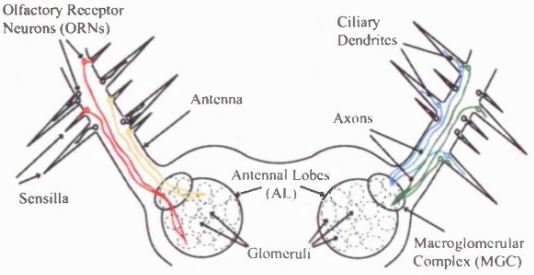

Figure 1.2: Basic layout of the moth brain. Two antennae protrude from the front of the head. The surface of these antennae is covered with thousands of tiny hair-like structures called sensilla. Each sensillum houses one or more ORNs whose main dendrites extend into the sensillum, and whose axons extend into the ALs. Axons from each type of ORN converge onto specific glomeruli. ORNs tuned to plant odours converge onto the main area of the AL, while ORNs tuned to sex pheromones converge onto glomeruli in the MGC.

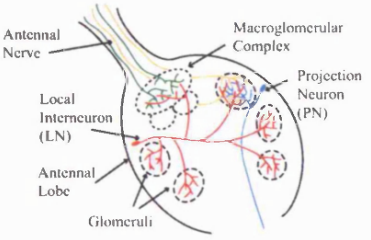

Figure 1.3: Antennal lobe neurons. The spherical regions are areas of densely packed neuropil known as glomeruli. The axons from a single type of ORN usually converge on the same glomerulus, and different ORN types converge on different glomeruli. LNs innervate many glomeruli. Dendrites from PNs can innervate a single glomerulus, or several glomeruli depending on their type, and their axons extend out to higher brain areas. (After Chong, 2003).

1.2.3 The Macroglomerular Complex

Although males and females usually have the same number of ordinary glomeruli in the AL, there is a degree of sexual dimorphism (Anton and Homberg, 1999; Christensen et al., 1989; King et al., 2000). Males have a small number of large glomeruli at the top of the AL in an area known as the macroglomerular complex (MGC), at which the axons from ORNs responsive to sex pheromones converge (for review see Boeckh et al., 1984). In the moth *Spodoptera littoralis*, the MGC consists of three glomeruli, a large glomerulus

atop two smaller glomeruli (Ochieng' et al., 1995). These glomeruli are known to respond to the major component of the sex pheromone, the minor component, and a behavioural antagonist (Ochieng' et al., 1995). The MGC is the main pheromone processing subsystem, and although some LNs innervate both glomeruli in the main AL and the MGC, on the whole there is very little interaction with the rest of the AL.

1.2.4 Neuron Types and Function

1.2.4.1 ORN Selectivity

ORNs respond most intensively to specific kinds of odour molecules, with other odour molecules producing varying response levels (Gustavsson et al., 1997). In this study, the ORNs of interest are those that detect sexual pheromones. Electrophysiological studies have shown that pheromone detecting receptor neurons on the male moth antennae are highly specific for single pheromone compounds (Hansson et al., 1989; Kaissling et al., 1989; Ljungberg et al., 1993). While these ORNs have been widely researched, work in this thesis appears to be the first to employ signal detection theoretic measures to assess the pheromone detection thresholds.

1.2.4.2 Local Interneurons

The LNs form a network linking together many glomeruli, with synapses onto both ORNs and PNs (see Figure 5.8 for examples). A large proportion of LNs show GABA-like immunoreactivity, supporting the physiological findings that output from LNs is inhibitory (Christensen et al., 1993; Distler, 1989). The number of glomeruli innervated by an LN varies greatly between neurons, and can range from LNs with homogeneous arborisations innervating almost all glomeruli (Anton and Hansson, 1994), to those known as oligoglomerular, innervating only a few (Anton and Hansson, 1994).

1.2.4.3 Projection Neurons

The PNs are commonly referred to as the output neurons of the AL, transmitting the olfactory information to higher brain functions for further processing. These PNs often innervate only a single glomerulus, although some have been shown to innervate several (Anton and Homberg, 1999 and references therein). PNs are greatly influenced by the inhibition they receive from LNs, but they also receive exclusively excitatory input from ORNs (Christensen et al., 1993; Hildebrand, 1996; Homberg et al., 1989; Sun et al., 1997). The complexity of the neuronal circuitry between ORNs, LNs and PNs within the AL glomeruli is thought to play an important role not only in the identification of odours, but also in the detection of the appropriate blend (Hansson and Anton, 2000 and references therein).

1.2.5 Synaptic Connectivity

Within this thesis, the responses of both ORNs and ALNs to olfactory stimulation are of interest, as is the synaptic connectivity within the glomerular structures. Glomeruli are densely packed structures containing a multitude of synaptic connections between ORNs, LNs and PNs. The incoming message is transferred synaptically from ORNs to LNs and PNs (Distler and Boeckh, 1996, 1997a,b). Since LNs often show projections to a large number of glomeruli (Anton and Hansson, 1994), specific innervation patterns are unlikely to provide insight into exact coding mechanisms, but are likely to distribute incoming signals across the AL. PNs on the other hand, often show uniglomerular innervation patterns (Anton and Homberg, 1999), and while little is known about the specific synaptic connectivity patterns of individual PNs, this synaptic circuitry provides the basis for the modelling study performed in this thesis.

1.2.6 The Olfactory Code

Information about the intensity and identity of an odour stimulus is encoded at the periphery in the pattern of activation of olfactory receptors, and is translated into action potential firing. The firing rates of ORNs are much faster than the rate of any modulation of the stimulus (odour plumes often contain pockets of pheromone arriving at a point of interception at frequencies around 4-10 Hz; Justus et al., 2002a; Murlis and Jones, 1981) suggesting the probability that the average rate of firing carries most of the stimulus information. It is possible however, that temporal factors may form part of the coding, and behavioural studies have shown that many moth species are unable to locate a source of odour unless the stimulus presentation is intermittent (Baker et al., 1985; Justus et al., 2002b; Mafra-Neto and Cardé, 1994, 1995b). This translates into two possible hypotheses, firstly, the ORNs could be encoding the stimulus in the form of precise patterns of spikes (within an individual ORN response), or secondly, the ORNs could be encoding the stimulus frequency dynamics in terms of variations in the firing rates or spike times across multiple neurons. While the actual solution is as yet unconfirmed, both of these ideas are investigated further in this thesis.

1.2.7 Further Information

This introductory chapter is intended to provide an outline of the basic structure of the olfactory pathway, to the level required for understanding the work in this thesis. Later chapters exist as self-contained units, and as such contain introductory material relevant to that specific area of work.

1.3 Outline of the Thesis

Chapter 2: General Techniques. A description of the recording techniques used by collaborators to obtain the electrophysiological data. The methods then used to extract the

spike timing information from the two types of neurons are explained, followed by a brief look at the verification of spike times.

Chapter 3: Sensitivity Enhancement in the Olfactory Pathway. An investigation into the detection thresholds for neurons at the first two stages of the olfactory pathway in response to a pheromone stimulus was performed. The thresholds found are compared to a theoretical model of signal convergence between the sensilla and the antennal lobe. A simple computational model is developed and used to demonstrate that the boost observed in the biology cannot be explained by simple pooling of information. Possible mechanisms that could explain the boost are discussed briefly, with further work in later chapters.

Chapter 4: Temporal Encoding at the Periphery. The calculation of detection thresholds in Chapter 3 assumes that the neurons in the sensilla use a change in the firing rate to convey the presence of the stimulus. This chapter looks at a couple of methods to investigate alternative encoding schemes at the periphery. Unitary Events Analysis (UEA) is then used to look for temporal structure in the ORN recordings.

Chapter 5: Modelling of Glomerular Mechanisms. Here, the modelling work is taken further with a look at the possible synaptic connections within a single glomerulus, and computations that individual neurons may be able to perform. A simple but biologically plausible computational model is created that demonstrates that when the glomerulus is considered as an active site of convergence (i.e. synaptic connections are capable of transforming input signals in some way), the signal can be boosted, thus enhancing the detection capability of an output neuron.

Chapter 6: Frequency Analysis of Olfactory Receptor Neuron Responses. Since it is known that the male moth must find a female by following an intermittent pheromone plume, this chapter looks at the frequency components found in the stimulus response of several neurons in the sensilla. The results are comparable with frequencies found in a wind tunnel generated plume, suggesting that the dynamics found in the neuronal responses cover the frequencies found in natural plumes.

Chapter 7: Conclusions.

Chapter 2

General Techniques

2.1 Chapter Overview

This chapter starts by describing the electrophysiological techniques used by collaborators at the Sveriges Lantbruksuniversitet in Sweden, which form the basis for experimental work related to this thesis. Recordings were taken from ORNs in the long sensilla on the antennae, and from neurons in the AL. The majority of neurons in the AL belong to one of two types, LNs or PNs. Discrimination between these two types of neurons is possible using morphological staining techniques, but these are not available for this data set. The term antennal lobe neuron (ALN) is used throughout the rest of this thesis to represent both neuron types.

The recordings obtained consist of a time series of amplified cell voltage values, sampled at 13888 Hz and 6380 Hz for ORNs and ALNs respectively. The two different methods used to take recordings have different spike sorting problems associated with them. For the method used to record from neurons within the AL, spike time extraction can be achieved by a simple threshold idea. The method used to record from the ORNs means that extracellular signals from multiple neurons may be present within the recording, and more complicated spike sorting methods may need to be employed.

Both methods of spike timing extraction, along with the necessary pre-processing performed on recordings are explained in this chapter. In the case of more than one neuron

picked up by the recording electrode, a method to determine which neuron is responding to the stimulus is described and verification techniques are briefly explained. The spikes used in subsequent analyses in this thesis were all sorted using the techniques described in this chapter.

2.2 Electrophysiological Recording Techniques

2.2.1 Insects

Experiments were performed on 1-5 days post-emergence male moths of *Spodoptera littoralis*. The moths had been reared for several generations on a potato-based diet (Hinks and Byers, 1976). The pupae were separated according to sex and kept in plastic boxes at 70% relative humidity, 23°C and a 16 h:8 h light/dark cycle. Adult moths were given excess of water until the start of the experiment.

2.2.2 Single Sensillum Preparation and Recording

Takuma Takanashi and Niels Skals took the single sensillum recordings used throughout this thesis.

The moth was mounted in a truncated pipette tip with one antenna protruding from the narrow end. The position of the antenna was fixed with dental wax (Surgident, Heraeus). To establish contact with the ORNs, tungsten microelectrodes that had been electrolytically sharpened in KNO_2-solution (Hubel, 1957) were used. The recording electrode and indifferent electrode (inserted into the abdomen) were positioned using a preparation microscope and micromanipulators. The signal was amplified using a high impedance amplifier (Syntech INR-02, Hilversum, The Netherlands). Signals from contacted ORNs were displayed on an oscilloscope (Tekscope, Tektronix), linked to a loudspeaker for audio monitoring. Recordings of signals were digitised using AutoSpike16 (Syntech, Hilversum, The Netherlands).

2.2.3 Intracellular Preparation and Recording

Qian Han took the intracellular recordings used throughout this thesis, with a number of preliminary recordings (used mainly for coding of spike sorting algorithms) taken by Anna-Carin Bäckmann.

The moth was restrained in a disposable plastic pipette tip, which was cut to allow passage of the head. The head position was fixed with dental wax (Surgident, Heraeus). The brain was uncovered and the sheath overlaying the AL was carefully removed. During the experiment, the brain was super-fused with a saline solution containing 150 mM NaCl, 3 mM $CaCl_2$, 3 mM KCl, 10 mM TES (N-Tris-methyl-2-aminoethanesulfonic acid) buffer, and 25 mM sucrose (pH 6.9) to increase osmolarity and prevent swelling of the brain tissue (Anton and Hansson, 1995).

Standard intracellular recording techniques were used (Kanzaki et al., 1989). The penetration sites were based on AL maps and previous results (glomerular representations (Carlsson et al., 2002; Sadek et al., 2002) and calcium imaging of responses (Carlsson et al., 2002)). Signals were displayed on an oscilloscope and recorded on FM tape. Signals were then digitised using AutoSpike32 (Syntech, Hilversum, The Netherlands).

2.2.4 Stimulation

A steady stream of charcoal-filtered and humidified air ($\sim$0.5 m s^{-1}) was continuously delivered through a glass tube ventilating the antenna (ipsilateral to the recording site for intracellular recordings, and the available one for single sensillum recordings). Ten microliters of solvent containing the given amount of stimulus was applied to a filter paper (5x15 mm) in a Pasteur pipette. The tip of this pipette was inserted into a small opening in the glass tube 20 cm from the antenna. A 4 ml s^{-1} 0.5 s air pulse was sent through the Pasteur pipette by means of a stimulation device (Syntech, Hilversum, The Netherlands). Various stimulus loads of the major component of the female-emitted sex pheromone, (Z, E)-9, 11-tetradecadienyl acetate (Z9, E11-14:OAc) were tested with inter-stimulus-intervals *ca.* 20 seconds, or until the neuron under study had fully recovered.

2.3 Extraction of Spike Timing

2.3.1 Extraction of ALN Spikes

Recordings from ALNs often showed fluctuations in the baseline voltage of the neuron, usually after the onset of the stimulus presentation (see Figure 2.1 a)). Although these fluctuations may contain important information relevant to the stimulus, this work is only interested in the spikes. Despite these fluctuations, most of the spikes have similar relative amplitudes, suggesting a thresholding method for finding the spike times. A profile plot from a recording (see Figure 2.1 c)) shows the need to first remove the baseline fluctuations by filtering the recording.

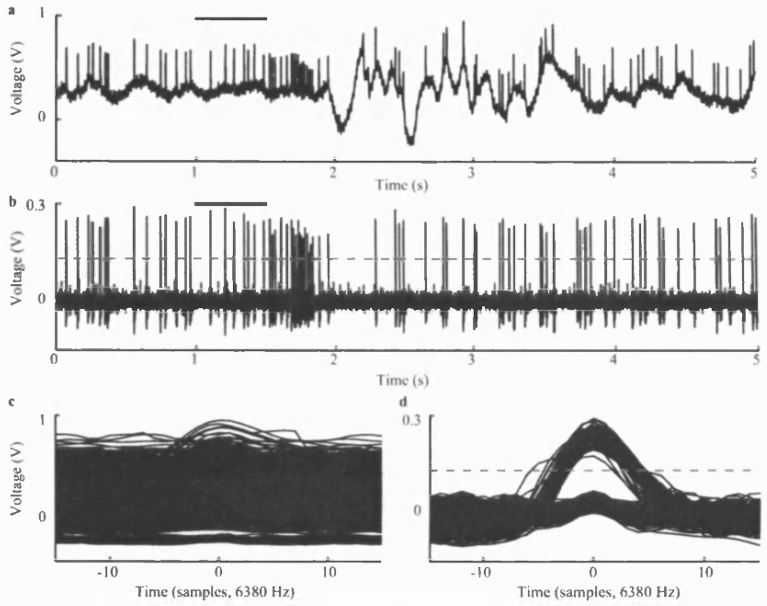

Figure 2.1: Spike Timing Extraction from ALNs. a) Raw Data from an ALN. As well as an increase in firing rate in response to the stimulus, recordings often showed an increase in the fluctuations of the baseline voltage of the neuron. Horizontal black bar represents the stimulus presentation. b) Filtered Data from an ALN. After filtering, the baseline fluctuations have been removed, and a threshold voltage can be set to identify spikes (red dashed line). c) Raw Profile Plot. Potential spikes (excursions of the voltage above the mean voltage level) are plotted, with the maximum value lined up with the 0 time point. The baseline fluctuations mean that the spike shapes and baseline cover the same voltage range and are indistinguishable. d) Filtered Profile Plot. After filtering, the baseline and spike shapes separate, and a threshold voltage can be set (red dashed line) with areas above this threshold representing spikes.

2.3.1.1 Filtering Procedure

A fourth order band-pass Chebyshev type 1 filter was designed to filter the recordings. The lower cut-off frequency was set to 100 Hz to ensure low frequency noise was removed (50 Hz power supply and the baseline fluctuations). The upper cut-off frequency was set to 1000 Hz to remove any high frequency noise that may exist in the system, while retaining as much spike information as possible (spike frequency ~320 Hz).

The recordings were processed twice, first the recording were filtered in the forward direction, and then the filtered sequence was reversed and run back through the filter. This ensured that the resulting sequence had zero phase distortion and had double the filter order. The resulting sequences had reduced amplitude, but the spike timing information remained the same, and the baseline fluctuations were removed (see Figure 2.1 b)). When a profile plot was created for the filtered recording, the spike shapes became visible (see Figure 2.1 d)), with good separation between spike shapes and background fluctuations (noise).

2.3.1.2 Threshold Method

With the baseline fluctuations removed, all the spikes had similar peak amplitudes. A threshold amplitude (usually set to be the mean plus four times the standard deviation of the filtered voltages) was selected, which varied for each recording, and sections of the time series above this threshold (red line in Figure 2.1 b) and d)) were deemed to be potential spikes. The actual spike time was taken to be the time at which the maximum voltage occurred.

2.3.2 Extraction of ORN Spike Timing

Recordings from ORNs in the long sensilla tended to have higher amplitude background noise in the baseline than those from ALNs (see Figure 2.2 for example), and since the recordings were extracellular, spikes from more than one neuron may be present (Ljungberg et al., 1993; Ochieng' et al., 1995; see Figure 2.2 b)). This posed a problem with using

the threshold method of spike sorting since an appropriate threshold value was not always obvious. However, the ORN spike shapes were consistent, and had a characteristic shape not typically seen in the baseline noise. This allowed the use of Principal Components Analysis (PCA; Appendix A) to determine which of a set of consecutive samples gave the most variation between the shapes of the spikes and the noise.

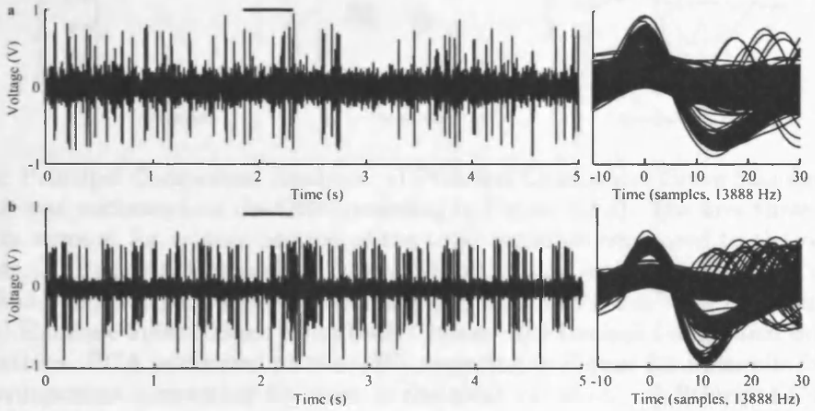

Figure 2.2: ORN Recordings. a) Raw Data from an ORN. Although each spike has a characteristic shape, in the profile plot, the spike shapes and baseline voltage shapes merge together, and a threshold method of spike sorting is not possible. b) Raw Data from a Recording with Two ORNs. Two spike shapes are seen in some recordings. In this case, the positive spike sections have similar amplitudes, but the negative sections show two different amplitudes. In the profile plot there are two groups of spike shapes corresponding to the two different neurons.

2.3.2.1 Principal Component Analysis Method (Lewicki, 1998)

To perform PCA, a list of potential spikes was generated for each recording, and aligned with sample zero being that of maximum amplitude across a series of samples (usually 10 samples before, and 30 samples after the sample with maximum amplitude, see Figure 2.2). PCA was performed on these sequences of samples using MATLAB (The Mathworks), and the standard deviations of principal component scores were plotted, showing that most of the variation occurred in the first 2-3 principal components (see Figure 2.3 a) and d)). When plotted against each other, the first two principal component scores usually gave groups of points which could be identified as corresponding to the different spike and noise shapes (see Figure 2.3 b) and e)).

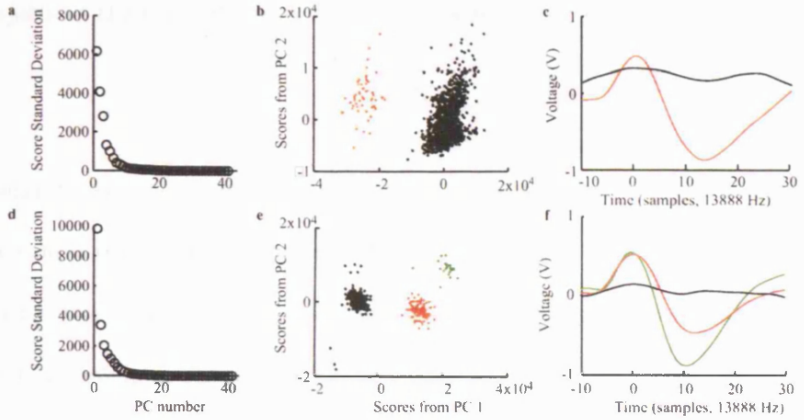

Figure 2.3: Principal Component Analysis. a) Principal Component Score Standard Deviations. PCA was performed on the ORN recording in Figure 2.2 a). The first three principal components account for a large portion of the total variation compared to the rest of the components. b) Principal Component Scores. When plotted against each other, the scores from the first two principal components form two groups, which can then be separated into clusters. c) Example Spike Shapes From Each Cluster. d) Principal Component Score Standard Deviations. PCA performed on the ORN recording in Figure 2.2 b) results in the first principal component accounting for most of the total variation. e) Principal Component Scores. When plotted against each other, the first two principal component scores form several possible groups, on which cluster analysis can then be performed. f) Example Spike Shapes From Each Cluster.

2.3.2.2 Clustering

The groups of points representing the different spike and noise shapes were separated into clusters using a hierarchical tree method. First, the Euclidean distance between pairs of points was computed using MATLAB's `pdist` function, where the distance, d_{rs}, between any two vectors, $\mathbf{x}_r$ and $\mathbf{x}_s$ is defined as

$$d_{rs} = \sqrt{(\mathbf{x}_r - \mathbf{x}_s)(\mathbf{x}_r - \mathbf{x}_s)^{\mathrm{T}}} \tag{2.1}$$

These pairs were then linked together giving successively smaller numbers of clusters. In general the cluster representing the baseline noise shapes contained far more points than clusters representing spike shapes, so the weighted pair-group method using arithmetic averages (WPGMA; Sneath and Sokal, 1973) was used with MATLAB's `linkage` function. Here, the distance between two clusters is calculated as the average distance between all

pairs of objects in the two different clusters, weighted by the size of the respective clusters

$$d(r,s) = \frac{1}{n_r n_s} \sum_{i=1}^{n_r} \sum_{j=1}^{n_s} dist\left(x_{ri}, x_{sj}\right) \tag{2.2}$$

where n_r and n_r are the numbers of objects in clusters r and s, x_{ri} is the ith object in cluster r and *dist* is the distance measure from Equation 2.1. The hierarchical cluster tree (Figure 2.4 a)) was then split into clusters, with the initial approach being to create more clusters than actual spike shapes, and then to group the excess clusters. Here, MATLAB's `cluster` function was used, which effectively draws an imaginary horizontal line across the hierarchical cluster tree such that the number of vertical lines bisected is equal to the requested number of clusters (red dashed line in Figure 2.4 a) with 4 clusters). The number of clusters specified were 3 for a neuron with single spike shape, and 4 for a neuron with two spike shapes (Figure 2.4 b)). To confirm which clusters were spike shapes, and which were noise, members of each cluster were plotted with the original recording data. Clusters with points lining up with the recording spike shapes corresponded to spikes from neurons, and all other clusters were noise, and were discarded (Figure 2.4 c)).

2.3.2.3 Multiple Neurons

Most of the recordings processed contained only a single spike shape, but in recordings where more than one neuron was present, it was important to determine which neuron was the one responding maximally to the stimulus. A simple method to determine this is to first sort the spikes as mentioned above, then plot the peri-stimulus time histograms (PSTHs) for each neuron, at several stimulus loads. Depending on the receptor types, several possibilities may occur:

- One neuron shows an elevated firing rate in response to the stimulus presentation, while other neurons show no response. In this case, the neuron that responds is the neuron of interest.

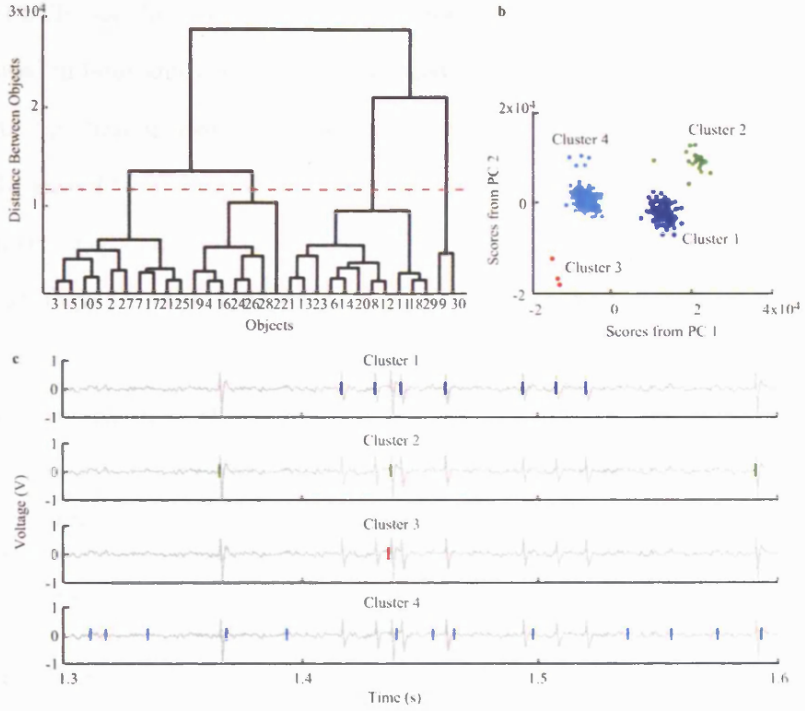

Figure 2.4: Cluster Analysis. a) Hierarchical Cluster Tree. Scores from the first two principal components of the recording in Figure 2.2 b) were first linked together in pairs according to their distance from other points. These small groups were progressively linked together with the weighted pair-group method using arithmetic averages, until all points were linked. The connections between the groups (starting when 30 groups remained) are shown, with a possible cut-off threshold (red dashed line) resulting in four clusters. b) Principal Component Score Clusters. The four clusters from a) are plotted, and form tight groups with a small number of stray points assigned to the closest clusters. c) Possible Spikes. The four recording segments show the locations of possible spikes from points within the four clusters. Points from Cluster 1 correspond to spikes with small negative amplitudes, Cluster 2 to spikes with larger negative amplitudes, and Clusters 3 and 4 don't line up with spikes in the raw data so correspond to the baseline noise shapes.

- All neurons show an elevated firing rate, but one neuron has a greater increase. In this case, the neuron with the largest increase is the neuron of interest.
- All neurons show an equal increase in the firing rate. This suggests that none of the neurons are maximally responsive to the stimulus, so the data should not be used.
- No neurons show an increase in the firing rate, so the data should not be used.

For the multiple neurons case shown in Figures 2.2 b), 2.3 and 2.4, clusters 1 and 2 are spike shapes from two different neurons, and clusters 3 and 4 are shapes from the baseline

noise. The PSTHs for these two neurons can be used to aid in the decision of which neuron should be used in later analyses. The spikes found from cluster 1 have an increase in firing rate to a 10^{-5} g stimulus load but little change in firing rate to other stimulus loads and the blank (Figure 2.5 a)), suggesting a response to the stimulus. The spikes found from cluster 2 show a slight increase in response to the stimulus presentation, regardless of the stimulus load (Figure 2.5 b)), suggesting a possible mechano-sensory response to the change in airflow caused by the stimulus presentation. For these recordings, only the neuron with spike shapes represented by cluster 1 would be used.

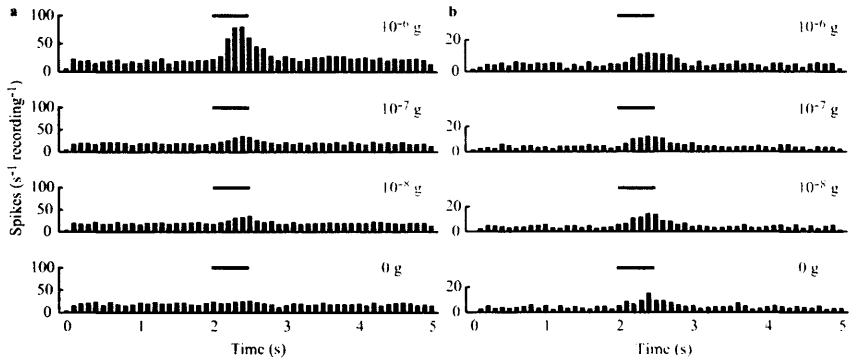

Figure 2.5: Choice of Responding Neuron. a) Peri-Stimulus Time Histogram for Cluster 1. Using time bins of 0.1 seconds, the PSTH for the 10^{-5} g stimulus load shows an increase in firing rate in response to the stimulus presentation (blank horizontal bars) whereas other stimulus loads and the blank show little change. b) Peri-Stimulus Time Histogram for Cluster 2. Using time bins of 0.1 seconds, the PSTH for all stimulus loads and the blank show a slight increase in firing rate due to the stimulus presentation, suggesting a mechanosensory rather than chemosensory response.

2.4 Verification of Spike Timings

As a simple measure to ensure that spikes were correctly identified, the sorted spike timings were superimposed onto plots of the raw data. Any misclassified spikes would be seen as sorted spike times with no corresponding raw data spike, or vice versa (Figure 2.6). A less time-consuming method is to look at the interspike intervals (ISI's) for each recording. All neurons have an absolute refractory period after a spike in which no other spikes can be generated. For ORNs, this period is $\sim$1 ms, so recordings with any ISI's less than 1 ms may have been clustered incorrectly (Figure 2.7).

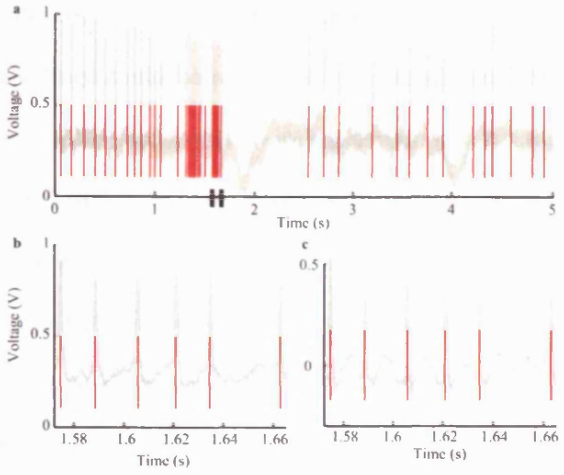

Figure 2.6: Spike Timing Visual Verification. a) ALN Raw Data and Spike Times. Spikes (red vertical lines) are superimposed on the raw data (grey trace) from an ALN. Misclassified spikes would be seen as either red vertical lines with no corresponding data spike, or a data spike with no corresponding red vertical line. b) and c) Spike Times with Original and Filtered Recordings. Zoomed in section corresponds to the portion of recording between the vertical black lines on the x axis in a). Spike times correspond to raw data spike both before and after the filtering process.

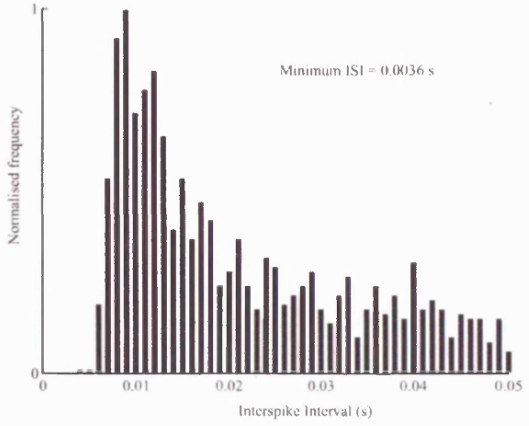

Figure 2.7: Inter Spike Interval Histogram. Time bins are 1 ms, with x axis truncated to 0.05 seconds. The absolute refractory period of a neuron means that there should be no interspike intervals of less than 1 ms. The histogram contains no intervals in the first three bins (i.e. < 3 ms) and in fact, the shortest interval is 3.6 ms.

2.5 Summary

Collaborators in Sweden took electrophysiological recordings from two areas of the olfactory system of the male moth *Spodoptera littoralis*, the antenna and the AL. Recordings from the

neurons in the AL contained fluctuations in the baseline voltage that could be filtered out such that a threshold voltage could be set to identify the spikes. Recordings from neurons in the antenna contained spike shapes that were consistent, but a threshold method of spike sorting was not possible. Instead, PCA was used, which formed groups of spike and baseline noise shapes when the scores from the first two principal components were plotted against each other. A weighted method of hierarchical cluster analysis was used to separate clusters and identify those clusters representing spike shapes. A simple method of comparing PSTH shapes over several stimulus loads was used to determine which the neuron of interest was (i.e. the neuron showing maximal response) if spikes from multiple neurons were present. Finally, spike times were verified visually, and ISI histograms were used to add confidence to the spike sorting and clustering techniques. The spikes sorted using the methods described in this chapter were used in the analyses of later chapters.

Chapter 3

Sensitivity Enhancement in the Olfactory Pathway

3.1 Chapter Overview

The main aim of this chapter is to quantify the detection thresholds of the first two stages of olfactory processing in the moth, ORNs on the antenna, and ALNs in the AL. These detection threshold estimates allow a direct quantification of the boost in the sensitivity achieved between these two processing stages in this animal. This boost is later compared to the boost expected due to the structure of the biological system.

Electrophysiological recordings were taken from these two types of neuron, and spike train timings were extracted as detailed in Chapter 2. Detection thresholds are here calculated using statistical methods and a detailed signal detection theoretic approach. The boost in the signal detection ability is defined as the difference in detection thresholds between the two pheromone processing stages.

A simple theoretical model is developed here. This model is based on the convergence of many signals onto a single integration site, and gives rise to an estimate of the boost in signal detection when Poisson spiking statistics are employed. When compared to this theoretical boost, the biological pheromone detection system demonstrates a higher than expected boost, well beyond that explainable by pooling of elevated firing rates, implicating alternative coding and readout mechanisms.

Assumptions made throughout this investigation are discussed, with none posing a significant effect on the outcome. Possible ways in which the biological system achieves this observed boost in signal detection are discussed, leading on to work performed to investigate these hypotheses in later chapters.

3.2 Introduction

The detection performance of a sensory system derives not only from the sensitivity of its underlying receptors, but also from the subsequent neural processing. In many sensory systems this processing achieves a boost in sensory signal-to-noise ratio, resulting in detection limits that are lower than can be supported by an individual receptor (Klein and Levi, 1985; Ulanovsky et al., 2003; Zanker and Harris, 2002). While many of the neural coding principles and processing mechanisms giving rise to these different forms of hyperacuity are unclear, one widespread strategy employed by the nervous system is the pooling of sensory signals from populations of receptor neurons to decrease uncertainty about the stimulus (Pearce et al., 2001a,b; Rieke et al., 1997). This arrangement does not preclude the possibility, however, that other neural coding strategies might be operating in tandem - for instance based upon the principles of temporal coding (Buonomano and Merzenich, 1995), synchrony (Gray et al., 1989; MacLeod et al., 1998) or noise-shaping (Mar et al., 1999).

Signal boosting strategies are likely to result from incremental and selective adaptations in sensory systems that require high levels of sensitivity to certain environmental stimuli (Maynard-Smith, 1998). The olfactory pathway of the moth is such an example, boasting extreme levels of sensitivity to highly specific pheromone stimuli. Minute quantities of pheromone at less than 10^{-15} M (Kaissling, 1971; Kaissling and Priesner, 1970), delivered as a blend (Tamaki and Yushima, 1974), are sufficient to elicit stereotypical behavioural responses essential for reproduction in this animal (Schneiderman et al., 1986), with as few as 6 pheromone molecules hitting the antenna sufficient to alter the heartbeat rhythm of the

male *Spodoptera littoralis* (Angioy et al., 2003). Reported detection thresholds for individual ORNs to their cognate pheromone ligands are too high to fully account for the detection capability necessary to support these behaviours (Mankin et al., 1991; Valeur et al., 2000), suggesting that sensory signal enhancement must be achieved through subsequent neuronal processing within the olfactory pathway. By acting as key points of convergence and thus sensory integration, glomeruli are likely to be important in this process, potentially lowering detection thresholds in different species (Boeckh and Ernst, 1987; Carlsson and Hansson, 2002; Duchamp-Viret et al., 1989; Hartlieb et al., 1997), by pooling sensory input from massive numbers of ORNs expressing identical olfactory receptors (Chen and Shepherd, 2002).

In order to assess the sensitivity boost possible during early olfactory processing, the detection thresholds to a major pheromone component (Z9, E11-14:OAc) at the first two stages of olfactory processing in an olfactory specialist, the moth *Spodoptera littoralis*, were calculated. By collecting a large data set of olfactory receptor and antennal lobe neuron electrophysiological recordings, a detailed signal detection theoretic analysis of sensitivity enhancement in this system was performed, hitherto not characterised in any animal with any statistical power. These detection threshold estimates allow a direct quantification of the boost in sensitivity achieved between the first two stages of the olfactory system in this animal.

Can the sensitivity boost observed in this system simply result from the pooling of elevated firing rates at the periphery, or conversely, are more subtle stimulus coding strategies at play? Specifically, since the early olfactory pathway is highly organised topographically in this animal, it was possible to compare the observed sensitivity boost with that suggested by theoretical considerations of signal convergence. The results in this chapter demonstrate that the sensitivity boost achieved during early olfactory processing by this animal is well beyond that explainable by pooling of elevated firing rates, implicating alternative coding and readout mechanisms.

3.3 Data Analysis Methods

3.3.1 Electrophysiological Recordings

Recordings from both ORNs and ALNs of the male *Spodoptera littoralis* were taken by collaborators in Sweden (see Chapter 2) in order to investigate and compare the detection thresholds at these two processing stages. The stimulus used was a 0.5 second square pulse of the major component of the female sex pheromone, Z9, E11-14:OAc. Stimulus loads ranged from 10^{-5} to 10^{-9} g for ORNs and 10^{-6} to 10^{-10} g for ALNs. Spike timings were extracted from each recording according to the methods described in Chapter 2.

3.3.2 Net Spikes Calculation

In order to quantify the detection thresholds, a measure of the neuronal response to the stimulus must first be defined. For this chapter, the activity measure used is the net spikes per second, i.e. the number of spikes above the background firing rate that were emitted in response to the stimulus. This method allows neurons with different background firing rates, and different response durations (see Figure 6.3 for examples), to be compared.

The net spikes, N, are calculated using the formula $N = S_S - S_P$ where S_S is the number of spikes occurring in a given time period t post-stimulus and S_P is the number of spikes occurring in the same time period t pre-stimulus. Since each neuron can have a different time lag between the onset of the stimulus and the onset of a response, the time period t was calculated individually for each neuron. This calculation was performed by splitting the repeat recordings into time bins to give the spikes per second per recording for the highest stimulus load for each neuron. The time bin in which the maximum spikes occurred was taken as the start point. The start and end time of t were taken to be the first of three consecutive time bins where the firing rate dropped below a threshold value (set to the mean background firing rate plus three standard deviations), before and after the maximum time bin respectively (see Figure 3.1). The net spikes s^{-1} measure can then calculated by dividing N by t to allow for comparison across different neurons.

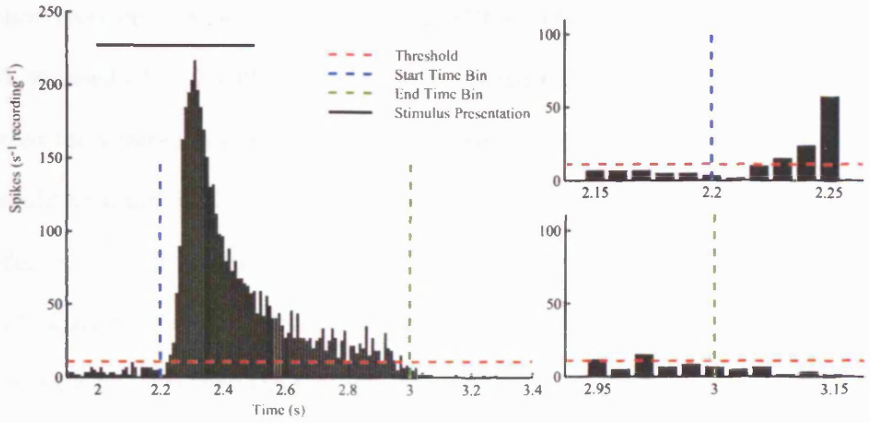

Figure 3.1: Calculation of Net Spikes Integration Period. The PSTH for a single ORN, at a stimulus load of 10^{-5} g, with a bin size of 10 ms is shown. The red-dashed horizontal line indicates the background firing rate plus three standard deviations. This is used as a threshold, where firing rates above this value are identified as the period in which the neuron is responding. The actual integration time is calculated by finding a start and end time bin. First the time bin containing the highest spikes s^{-1} was found. The start time bin was found by moving left along the time axis until three consecutive time bins contained fewer than the threshold number of spikes s^{-1} (blue dashed line, see also top right). The end time bin was found by moving right along the time axis until three consecutive bins contained fewer than the threshold number of spikes s^{-1} (green dashed line, see also bottom right). In this example, the net spikes integration period, t, is from 2.2 - 3 s.

3.3.3 Statistical Methods

To accurately define where the detection thresholds lie, multiple t-tests are used. Each individual t-test is used to determine whether the response to a single stimulus load is statistically different to the response to the blank stimulus, and thus detectable. In this case, the null hypothesis used is that the means of the net spikes s^{-1} with the stimulus load (μ_s) and blank (μ_b) are from the same parent distribution, $H_0 : \mu_s = \mu_b$, with the alternative hypothesis being that the presentation of a stimulus load elicits a larger response than the blank, $H_1 : \mu_s > \mu_b$.

Since multiple t-tests were used to compare each stimulus load with the blank stimulus presentation, the probability of making one or more type 1 errors (rejection of the null hypothesis when it is true) over the "family" of t-tests increases. An adjustment can be made such that a significance level, α^*, is used for each comparison, ensuring the overall

significance level of α is not exceeded (Hsu, 1996). For an overall significance level of $\alpha = 0.05$, we used $\alpha^* = 0.0102$ for each of our five individual comparisons. It is important to note that the t-tests must be independent in that the set of blank recordings used for each t-test should be a different set to those used for each of the other t-tests. To achieve this independence, for each stimulus load, 1000 t-tests were performed with randomly selected subsets of blank recordings (without replacement) for both ORNs and ALNs. The final p-value for each t-test is then the average of the p-values generated during the 1000 repeats.

3.3.4 Receiver Operating Characteristic (ROC) Curves

The receiver operating characteristic (ROC) curve, and corresponding area under the ROC curve (AUC), provide a non-parametric measure of a neurons ability to discriminate between a stimulus being present or not, and hence give an estimate of the detection thresholds. The ROC curve is a plot of the probability of a neuron responding to a stimulus load ("hits") against the probability of a neuron responding to the blank stimulus ("false alarms"; Green and Swets, 1966). A control curve was created to represent the case of identical distributions of stimulus present and stimulus not present. Recordings to the blank stimulus were randomly sampled (with replacement) and were assigned to one of two groups. One of these groups was treated as if a stimulus load was present, and the other was assumed to be the blank stimulus. The hits and false alarms were then calculated. The separation process was repeated a number of times (e.g. 30 or 50 times) and the mean values of the hits and false alarms were plotted against each other to form the control curve.

The AUC represents the probability of correctly identifying when the stimulus is present. An AUC close to 0.5 represents the chance level of completely overlapping stimulus present and stimulus not present distributions. Perfect discrimination is signified by an AUC of 1 and represents totally separate distributions. The AUC has been shown to be equivalent to the value of a Mann-Whitney or Wilcoxon non-parametric test (Bamber, 1975) and this relationship between the AUC and the Wilcoxon statistic has been used to derive

statistical properties such as the Standard Error (SE; Hanley and McNeil, 1982). From this relationship, a critical ratio, z, can be calculated (Hanley and McNeil, 1982, 1983) to determine whether the difference in the area under two ROC curves is due to random or real effects.

3.4 Results

3.4.1 Concentration-Specific Responses at Periphery and Antennal Lobe

3.4.1.1 Recorded Data Sets

Recordings from many ORNs at the periphery and ALNs in the AL were taken with the aim of assessing changes in response due to varying stimulus quantities. The stimulus used was the major component of the female sex pheromone, Z9, E11-14:OAc since this is vital for the sexual reproduction of *Spodoptera littoralis*, with neurons tuned to this component in plentiful supply. Overall, 28 ORNs were tested with stimulus loads from 10^{-9} g to 10^{-5} g (mean numbers of between 8 to 33 repeat recordings were taken at each stimulus load, but only the first 5 recordings were used in statistical analyses), and 33 ALNs with stimulus loads from 10^{-10} g to 10^{-6} g. Both ranges used stimulus loads increasing in decadic steps, and represented the range in which the detection thresholds were expected to lie (Anton and Hansson, 1995; Ljungberg et al., 1993). Neurons were also tested with blank stimulations (pipette and carrier liquid presented with no pheromone) to assess the proportion of the neuronal response that is due to airflow changes.

3.4.1.2 Repeatability of Responses

Individual neurons tuned to the same chemical may produce different patterns of firing in response to a stimulus, but recordings from a single neuron were found to be highly repeatable (see Figure 3.2). This is important for the analysis in this chapter since several recordings from each ORN, at each stimulus load, are used to assess the elevation in firing rate. Elevated and stereotypical firing rates are best seen at higher stimulus loads, where

the response stands out from the background firing. An example of an elevated firing rate response is shown in Figure 3.2 a). Repeat recordings using a stimulus load of 10^{-5} g are shown from the same ORN. Both recordings show a response characterised by a period of high firing activity during stimulus presentation followed by a long silent period (most likely to be caused by short-term adaptation). Moreover, this characteristic response pattern was conserved over all the recordings taken from this ORN, as demonstrated in the raster plot of Figure 3.2 b) and the peri-stimulus time histogram (PSTH) in Figure 3.2 c). It can be seen that ALNs also respond in a repeatable manner (see Figures 3.2 d)-f)) with elevated firing rates usually beginning after stimulus presentation followed by a weaker period of silence (caused by either adaptation, or inhibition from other neurons).

3.4.1.3 Concentration-Specific Responses

In order to understand how neuronal responses depend upon stimulus load, raster plots were constructed from recordings taken from both ORNs and ALNs (see Figure 3.3 a)). Most of the ORNs recorded show elevated firing rates in response to higher stimulus loads (subplots iv and v in Figure 3.3 a)). Whereas ALNs show a wider variety in background firing rates than do ORNs, they also show smaller increases in firing rate due to the stimulus (see Figure 3.3 b)).

As a more quantitative analysis of elevated firing rates in response to increasing stimulus load dose response curves for the ORNs and ALNs were constructed (see Figure 3.3 c) and d)), which allow the comparison of the response due to the stimulus across all the stimulus loads tested. Since the above background firing rates of the neurons in response to the stimulus are of interest, the net spikes s^{-1} (see Section 3.3.2) were calculated and used here. This allows the comparison of neurons with differing background firing rates by quantifying the response in terms of the number of excess spikes produced above the background rate. The individual points in Figures 3.3 c) and d) represent the net spikes s^{-1} for each recording at each stimulus load for the ORNs and ALNs respectively, with the solid lines showing the

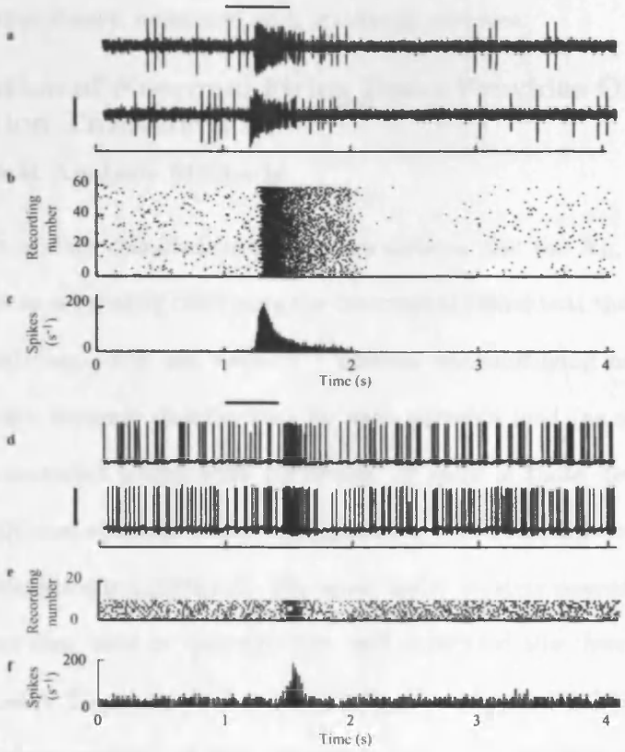

Figure 3.2: Individual Response Properties of Olfactory Receptor and Antennal Lobe Neurons to the Major Pheromone Component. a) Example ORN Recordings. Two example recordings taken from the same ORN to a stimulus load of 10^{-5} g (vertical scale bar: 1 Volt, horizontal bar represents stimulus duration). Notice the prolonged period of silence in the response after a period of strong bursting activity. b) ORN Raster Plot. Raster plot of many repeated presentations at a 10^{-5} g stimulus load from the same ORN as in a) (a blank presentation occurred between each stimulus presentation, data not shown). This neuron shows a highly repeatable response to the stimulus. c) ORN PSTH. PSTH for the spike trains in b), with a bin size of 10 ms. There is little spontaneous firing with a large increase in the firing rate due to the stimulus presentation. This elevated firing rate decreases over a period of ~1 s, followed by a period of inhibition, then returns to the background firing rate. d) Example ALN Recordings. This neuron shows a repeatable response to this stimulus. f) ALN Raster Plot. Raster plot of many repeated presentations at a 10^{-7} g stimulus load from the same ALN as in d). This neuron shows a repeatable response to this stimulus. e) ALN PSTH. PSTH for the spike trains in e), with a bin size of 10 ms. The neuron shows increased latency compared to the ORN with a short-lived increase in firing rate from the background rate due to the stimulus.

average dose response curves for the population of neurons. Higher stimulus loads for both ORNs (10^{-5} g, 10^{-6} g and 10^{-7} g) and ALNs (10^{-6} g and 10^{-7} g) generated more net spikes s^{-1} than blanks, suggesting a genuine response. The dose response curves do not in themselves provide an accurate estimate of detection thresholds, so these were assessed

using signal detection theory, combined with statistical analyses.

3.4.2 Evaluation of Neuronal Firing Rates Provides ORN and ALN Detection Thresholds

3.4.2.1 Statistical Analysis Methods

The large amounts of data collected from both the antenna and the AL, enable the use of statistical measures to accurately determine the detection thresholds at these first two stages of the olfactory pathway. The net spikes s^{-1} method of quantifying neuronal responses was used to generate separate distributions for each stimulus load (as seen in Figure 3.5 right). Multiple one-tailed t-tests were performed on pairs of these distributions (blank compared with individual stimulus loads, see Section 3.3.4) to determine whether differences in response were statistically significant. The areas under receiver operating characteristic (ROC) curves were also used to quantify how well separated the distributions were for varying stimulus loads. Together, these two methods, when applied to both ORN and ALN responses, give a robust estimate of their respective detection threshold, which provide one measure of the sensitivity enhancement between the first two stages of the olfactory pathway.

3.4.2.2 Assumption of a Homogeneous Population of ORNs

In the olfactory pathway of *Spodoptera littoralis*, odour molecules are detected by ORNs present in sensory hairs, sensilla, on the moth antenna. Approximately 10,000 sensilla containing Z9, E11-14:OAc-specific ORNs are present on each male antenna. Each sensillum contains two ORNs, identifiable by firing action potentials (spikes) of different amplitude. The large spike-firing ORN is tuned to the major component of the sex pheromone, Z9, E11-14:OAc, with a very high specificity. It has been suggested that each of these ORNs expresses a single type of olfactory receptor (Buck, 1996; Chess et al., 1994; Ngai et al., 1993). Given these ideas, all the ORNs recorded from are assumed to be a homogeneous population and to behave in a stereotypical manner when stimulated with a pheromone component. For the statistical analyses, all ORNs are considered as members of this population, although only the first five stimulus presentations to each ORN (at each stimulus load) were used to reduce possible effects of adaptation to the stimulus.

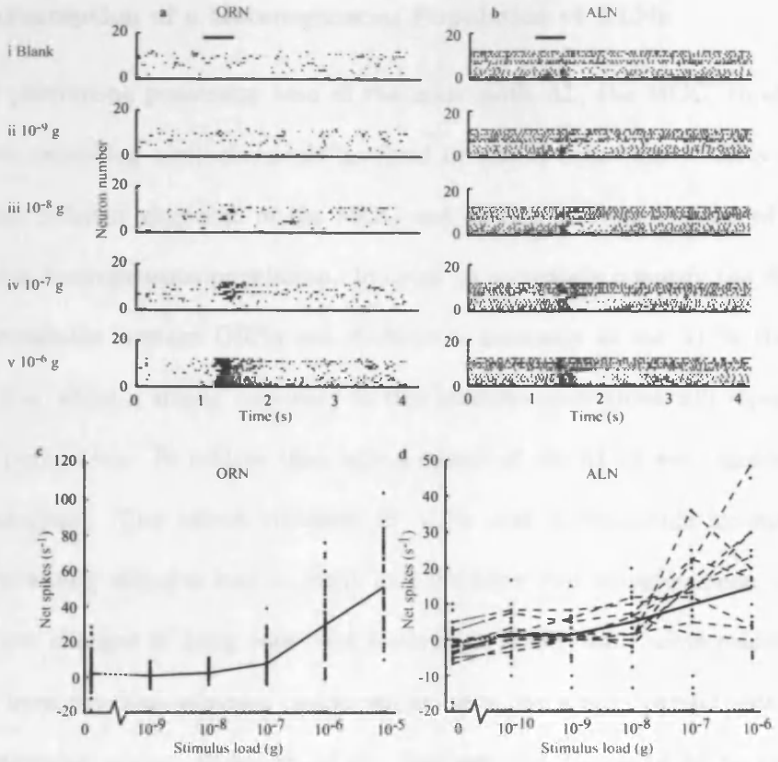

Figure 3.3: Stimulus Load Dependence of Olfactory Receptor and Antennal Lobe Neurons. a) ORN Raster Plots. Raster plots of the spikes from 12 different ORNs at each of 4 stimulus loads and a blank presentation. The order in which the neurons are plotted is preserved in each subplot. Most of the neurons show an elevated firing rate in response to the stimulus loads of 10^{-6} g and 10^{-7} g while responses to lower stimulus loads are not obvious. b) ALN Raster Plots. Raster plots of the spikes from 12 different ALNs at the same 4 stimulus loads and a blank presentation. The order in which the neurons are plotted is again preserved in each subplot. The difference in background firing rates across the neurons is evident with some neurons only firing in response to the stimulus. The firing rate of most neurons is elevated in response to the higher stimulus loads of 10^{-6} g and 10^{-7} g. Although some of these neurons show inhibition, all exhibit an initial period of excitation to the stimulus, and it is this period of excitation that is used to estimate the detection thresholds. c) ORN Dose Response Curves. Dose response curves for the ORNs using the net spikes method. Individual points show the net spikes for each of the recordings and the solid line shows the population average. The majority of ORN recordings, and the population average, have higher net spikes s^{-1} in response to stimulus loads of 10^{-5} g and 10^{-6} g compared with the blank stimulus. The recordings and population average net spikes s^{-1} in response to the lower stimulus loads are not noticeably different to the response of the blank stimulus. d) ALN Dose Response Curves. Dose response curves for the ORNs using the net spikes method. Individual points show the net spikes for each of the recordings and the solid line shows the population average. The dashed lines represent a subset of ALNs containing only neurons with a higher response, in terms of the number of net spikes s^{-1}, to all stimulus loads compared with the blank stimulus. The overall trend of the net spikes s^{-1}, for the population and individual recordings, is an increase in response to increasing stimulus loads.

3.4.2.3 Assumption of a Heterogeneous Population of ALNs

Within the pheromone processing area of the male moth AL, the MGC, three glomeruli receive input regarding semiochemicals involved in sexual behaviour. ALNs specifically innervate the different glomeruli of the MGC and display different degrees of specificity, representing a heterogeneous population. In order to accurately quantify the difference in detection thresholds between ORNs and ALNs it is necessary to use ALNs that respond maximally (i.e. show a strong response) to this stimulus since these will represent a homogeneous population. To achieve this, only a subset of the ALNs was considered in the statistical analyses. This subset consisted of ALNs with a monotonic increase in firing rate with increasing stimulus load at blank and the lower two stimulus loads. ALNs with non-monotonic changes in firing rate were assumed to either show some response to contamination from previous stimulus presentations, or to have only partial affinity for this particular stimulus odour. Examples of the dose response curves of ALNs used both in the subset, and left out of the subset are shown in Figure 3.4 a) and b) respectively. The complete subset is shown as dashed lines in Figure 3.3 d).

3.4.2.4 Application of t-tests to ORN Population and ALN Subset

Multiple one-tailed t-tests were performed on distributions of the net spikes s^{-1} measure across the ORN population (Table 3.1 a)). These data demonstrate that the net spikes s^{-1} generated in response to 10^{-7} g stimulus load and above derives from a statistically different parent distribution than to the blank stimulus, whereas lower stimulus loads are not statistically different. Thus, an estimate of the detection threshold for the population of ORNs can be found to be not more than 10^{-7} g for this pheromone component when presented alone.

Multiple one-tailed t-tests were performed on the ALN subset (Table 3.1 b)). The resultant p-values demonstrate that the net spikes s^{-1} measure obtained across all stimulus loads tested are from statistically different parent distributions to that of the blank stimulus.

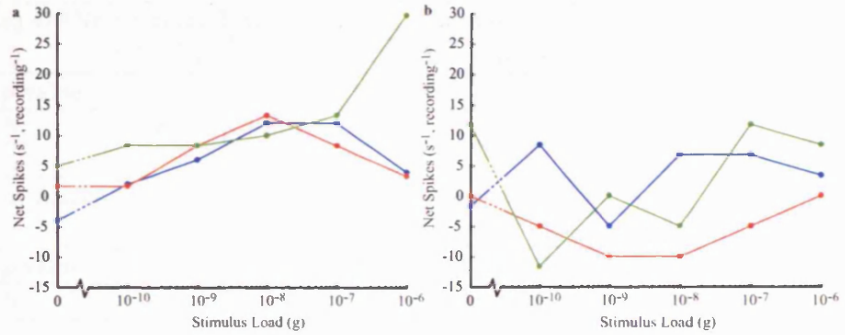

Figure 3.4: Selection of Quasi-Homogeneous ALN Subset from Heterogeneous ALN Population. a) Examples of ALNs Included in the Subset. A number of ALNs were selected from the whole population to represent a homogeneous subset. The requirement for the ALNs were that they had a monotonic increase in firing rate with increasing stimulus load for blank and the lower two stimulus loads, 10^{-10} g and 10^{-9} g (or no increase). For the three examples shown, the blue example shows monotonic increases between the blank, 10^{-10} g and 10^{-9} g stimulus loads. The red and green examples show monotonic increases between some stimulus loads, and no increases between others. b) Examples of ALNs Excluded from the Subset. These three examples of ALNs were part of the heterogeneous population, but not part of the homogeneous subset. Increases in stimulus load produce either an increase or a decrease in firing rate, some of which are not monotonic. These neurons do not appear to be maximally responsive to the stimulus load tested, so were not included in the statistical analyses.

The detection threshold for the ALNs can therefore be estimated as not more than 10^{-10} g for this pheromone component when presented alone. Note that this is a conservative estimate, however, since lower stimulus loads were not tested in this experiment and all stimulus loads tested produced statistically different responses compared to blank.

3.4.2.5 Application of ROC and AUC Analyses to ORN Population and ALN Subset

ROC curves were then generated (see Section 3.3.4) to assess the ability of the neurons to discriminate between a stimulus being present or not. The AUCs provide a non-parametric method of estimating the detection thresholds. For this analysis, the whole population of ORNs, and the same subset of ALNs were used. The distributions of net spikes s^{-1} used to create the ROC curves are shown on the right of Figures 3.5 a) and b).

For ORNs, the pairs of distributions are well separated for higher stimulus loads (10^{-5} g and 10^{-6} g), allowing for good discrimination between stimulus and blank (see Figure 3.5

a) ORNs	t-tests, blank compared to stimulus loads:					
	10^{-5} g	10^{-6} g	10^{-7} g	10^{-8} g	10^{-9} g	10^{-10} g
p-value	0*	0*	0.00998*	0.93	0.99802	-
N	202	213	208	202	203	-

b) ALNs	t-tests, blank compared to stimulus loads:					
	10^{-5} g	10^{-6} g	10^{-7} g	10^{-8} g	10^{-9} g	10^{-10} g
p-value	-	0.0004*	0.00042*	0.00333*	0.00617*	0.00527*
N	-	22	22	22	22	22

Table 3.1: Demonstration of Detection Thresholds for Olfactory Receptor and Antennal Lobe Neurons based on Second Order Statistics. Asterisks show statistically significant differences ($p < 0.0102$, multiple t-tests). a) ORNs. Multiple one-tailed t-tests comparing the net spikes s^{-1} produced by the population of ORNs in response to individual stimulus loads and the blank stimulus. The p-values suggest that the mean net spikes s^{-1} from the stimulus loads of 10^{-7} g and above are from statistically different distributions than the blank stimulus. This gives an estimate of the detection threshold as 10^{-7} g. b) ALNs. Multiple one-tailed t-tests comparing net spikes s^{-1} produced by the selected subset of ALNs in response to individual stimulus loads and the blank stimulus. The p-values suggest that all stimulus loads tested produce distributions of net spikes s^{-1} that are statistically different to the blank stimulus distribution. This gives an estimate of the detection threshold of $< 10^{-10}$ g.

a)). This high level of discrimination is demonstrated on the left plot by these stimulus loads having AUC values close to 1, implying a probability of misclassification of less than 10%. Distributions of net spikes s^{-1} for the lower stimulus loads show strongly overlapping distributions with that obtained from the response to blanks. This signifies poor detection accuracy in determining whether the stimulus is present or not, and is demonstrated by the ROC curves lying close to the chance line (dashed line along the diagonal), where the probability of error is close to 50%.

For ALNs, the distribution of net spikes s^{-1} at each stimulus load begin to overlap the blank distribution as the stimulus load decreases, but in all cases, there is a clear separation between the mean values of the distributions (see Figure 3.5 b)). This is represented by the ROC curves for all stimulus loads lying in the top left of the plot, resulting in detection performance well above chance. The AUC values of 0.81405 for the stimulus load of 10^{-10} g suggests that even at this low level, individual ALNs are able to correctly determine when

the stimulus is present 81% of the time. Comparisons of the AUC values for each stimulus load with blank, allow for the to calculation of the z-ratio (or z-score from a standard normal table) used to identify detection thresholds (see Section 3.3.4). Values of the z-ratio above 2.569 correspond to the p-value < 0.0102 used in the multiple t-tests.

The z-ratio values for the set of ORN recordings show that stimulus loads of 10^{-7} g and above produce statistically different ROC curves to that of the control curve, but lower stimulus loads do not (see Table 3.2 a), significantly different AUC values (z-ratio ≥ 2.569) denoted with an asterisk). This provides an estimate of the detection threshold as being 10^{-7} g, which corresponds to the threshold that was estimated using multiple one-tailed t-tests. In contrast, for the subset of ALNs tested, the z-ratios show that all the stimulus loads produce ROC curves that are statistically different from the ROC for blank (see Table 3.2 b)). This statistical analysis gives an estimate of the detection threshold as being $< 10^{-10}$ g, which again corresponds to the threshold estimated using multiple one-tailed t-tests.

a) ORNs	Control compared with stimulus loads:					
	10^{-5} g	10^{-6} g	10^{-7} g	10^{-8} g	10^{-9} g	10^{-10} g
AUC	0.99534	0.94185	0.68639	0.60955	0.50785	-
z-ratio	13.826*	10.909*	3.4323*	1.9183	0.03049	-

b) ALNs	Control compared to stimulus loads:					
	10^{-5} g	10^{-6} g	10^{-7} g	10^{-8} g	10^{-9} g	10^{-10} g
AUC	-	0.95868	0.97734	0.88017	0.8595	0.81405
z-ratio	-	3.7328*	3.9401*	3.5572*	3.4391*	3.2475*

Table 3.2: Determination of Detection Thresholds for Olfactory Receptor and Antennal Lobe Neurons based on Receiver Operating Characteristics. Asterisks show statistically significant differences (z-ratio≥ 2.569). a) ORNs. Comparison of the AUC of each stimulus load with the AUC of the control for the whole population of ORNs. The z-ratio suggests that stimulus loads of 10^{-7} g and higher produce ROC curves that are statistically different to the control curve. This gives an estimate of the detection threshold as 10^{-7} g. b) ALNs. Comparison of the AUC of each stimulus load with the AUC of the control for the selected subset of ALNs. The z-ratios suggest that all stimulus loads tested produce ROC curves that are statistically different to the control curve. This gives an estimate of the detection threshold as $< 10^{-10}$ g.

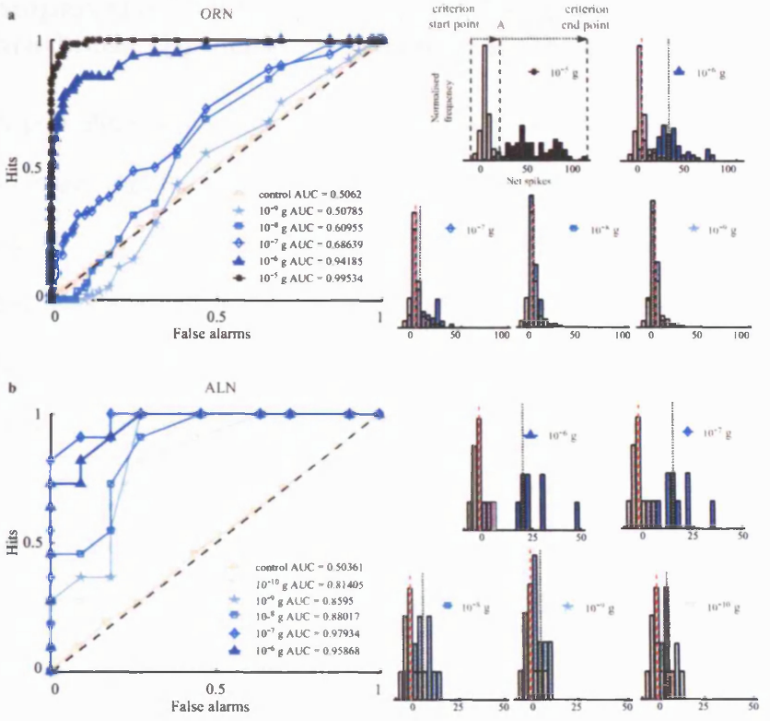

Figure 3.5: Receiver Operating Characteristic Curve Analysis of Olfactory Receptor and Antennal Lobe Neurons. The symbols and colours used for each stimulus load are conserved in both ROC figures and in the distribution figures. The AUCs are also shown, which is directly related to the psychophysical discriminability parameter d'. a) Left: ORN ROC Curves. ROC curves generated at each stimulus load for the ORNs. Stimulus loads of 10^{-5} g and 10^{-6} g have AUCs close to 1 suggesting almost perfect discrimination between stimulus present or not. Lower stimulus loads are closer to chance levels (dashed diagonal). a) Right: Net Spikes Distributions. Distributions of net spikes used to create the ROC curves. In each subplot, the net spikes distribution for the blank stimulus is depicted in pink. The mean net spikes s^{-1} for the blank distributions are indicated by vertical red dashed lines, and by vertical black dotted lines for the stimulus distributions. To create ROC curves, a criterion start point is set below the minimum number of net spikes s^{-1} in the distributions and incremented until higher than the maximum number of net spikes s^{-1}. At a given point A, blank net spikes above the criterion indicate 'false alarms' and stimulus net spikes above the criterion indicate 'hits'. The ROC curve is then a plot of the hits against false alarms. The distributions of stimulus and blank net spikes are well separated with stimulus loads of 10^{-5} g and 10^{-6} g, indicated by the separation between the mean values of the distributions, but these distributions overlap significantly as the stimulus load decreases. At stimulus loads of 10^{-8} g and 10^{-9} g the stimulus load distributions lie on top of that of the blank. b) Left: ALN ROC Curves. ROC curves generated at each stimulus load for the subset of ALNs. All stimulus loads have AUCs above 0.8, indicating a high level of discrimination over the range of stimulus loads tested. b) Right: Net Spikes Distributions. Distributions of net spikes used to create the ROC curves. In all cases, the distribution of net spikes for the stimulus loads lie to the right of the distribution of blank net spikes, as demonstrated by the separation of the mean values of the distributions, indicating a response to the stimulus.

3.4.3 Comparative Analysis Between ORN and ALN Detection Thresholds Quantifies Sensitivity Enhancement

Two different procedures were used to determine at what stimulus load the response from neurons is no longer significantly different to the response due to a blank stimulus. The first method used, the multiple t-test, is parametric in that it assumes the data under consideration are normally distributed. This method yields detection thresholds of 10^{-7} g and $< 10^{-10}$ g, for ORNs and ALNs, respectively. The second method used, the ROC and associated AUC analysis, provides a non-parametric method of determining the discriminability of a stimulus based on ranking the data. This method generates the same detection thresholds as the t-test method for both ORNs and ALNs for the pheromone component tested. It is possible then to conclude that there exists a reduction in detection threshold of at least 3 orders of magnitude between the neurons at the periphery and in the AL, which corresponds to an equivalent boost in sensitivity.

3.4.4 Finer Temporal Resolution Fails to Improve Discriminability in ORN Responses

While the signal detection theoretic analysis conducted for ORN responses employed rigorous criteria to specify over which time scale the spike events should be integrated, it was necessary to try to discount the possibility that the discriminability of the stimulus could be enhanced by considering alternative timescales. To do this, ROC curves were constructed for ORN responses over many independent time periods, of differing lengths, from the start of the stimulus presentation to the time at which the neurons appeared to stop responding. Integration periods of 50 ms, 100 ms, 200 ms and 500 ms were investigated, with each window moving along by 10 ms to form the next integration period. For each test period, the AUC values were compared against that of the control curve, and detection thresholds were calculated using the same criteria as previously employed (see Figure 3.6 a)-e) for example time periods and ROC curves). These detection thresholds, for each of 590 different integration time periods, are shown in Table 3.3. In all cases, the discrimination performance was

worse than that shown in the original analysis (see Figure 3.5 a)), leading to the conclusion that alternative integration times do nothing to improve stimulus discriminability.

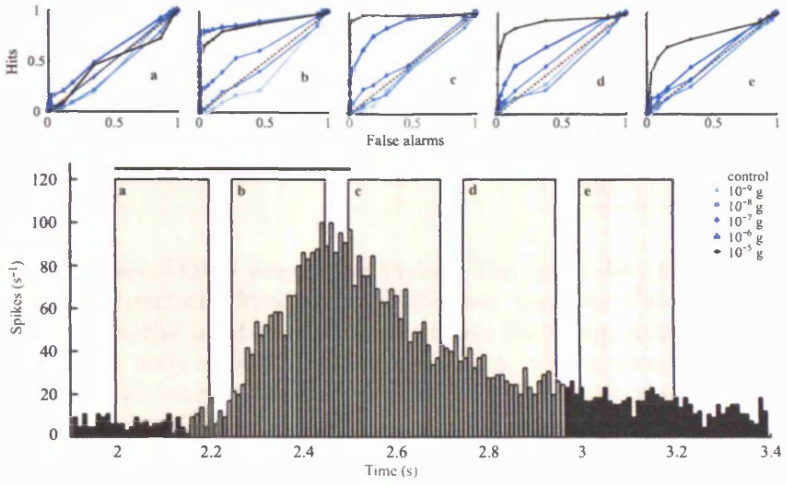

Figure 3.6: ROC Analysis of Net Spikes Integration Period. The symbols and colours used for each stimulus load in the ROC curves correspond those used in Figure 3.5. Lower Panel: PSTH For An ORN. Horizontal bar shows the stimulus duration. The PSTH for an ORN, at a stimulus load of 10^{-5} g is shown. Bars in grey correspond to the integration period used when determining the detection threshold for this ORN. Shaded areas labelled a) to e) show a subset of the integration periods tested. Upper Panel: ROC Curves For Labelled Integration Periods. a) After the stimulus has started, but before the ORN seems to respond, the ROC curves for all stimulus loads lie close to the chance line (dashed black line along the diagonal). b) At this point the ORN is now responding to the stimulus, and the ROC curves for the two highest stimulus loads have moved away from the chance line. c) The ORN is still responding to the stimulus with an elevated firing rate, and the curves from the two highest stimulus loads are further from the chance line. d) The response of the ORN is now dying down, corresponding to the ROC curves returning closer to the chance line. e) Only the firing rate of the ORN in response to the 10^{-5} g stimulus load is still slightly elevated, demonstrated by all other ROC curves again lying close to the chance line.

3.4.5 Inhibitory Responses of ALNs Show Greater Sensitivity to Pheromone Stimuli

3.4.5.1 Raw Inhibited ALN Data

The majority of ALNs encountered showed stereotypical responses: elevated firing rates during stimulus presentation, followed in fewer cases by a prolonged suppression in firing rate (Figure 3.3 b)). However, there were a small number of ALNs showing an inhibitory response to far lower stimulus loads. These particular ALN responses were studied at further stimulus loads to provide an estimate of their detection threshold. The response of one such

	Length of Time Period				
Threshold	50	100	200	500	Total
10^{-9} g	0	0	0	0	0
10^{-8} g	0	0	0	0	0
10^{-7} g	0	0	3	28	31
10^{-6} g	6	15	23	4	48
10^{-5} g	6	15	16	18	55
None	138	120	98	100	456
					590

Table 3.3: Investigation of ORN Integration Period. The signal detection theoretic method of determining the detection threshold of ORNs was used on a number of alternative timescales to discount the possibility of lower thresholds being attainable. Independent time periods, from the start of the stimulus presentation, moving along by 10 ms each time, to where the neurons no longer showed any response, with durations of 50 ms, 100 ms, 200 ms and 500 ms were used. The AUC values were compared with the control curve, and the detection threshold was set to be the lowest stimulus load at which z-ratio ≥ 2.569. The numbers of integration periods with thresholds at different stimulus loads are shown. Integration periods where even the highest stimulus load was not significantly different from the control are entered in the row labelled 'None'. From 590 different integration periods, only 31 (5.25 %) matched the detection threshold of 10^{-7} g found during the original analysis (Table 3.2), with no integration periods yielding a lower detection threshold.

neuron is shown in Figure 3.7 a). This neuron shows an inhibitory response to a stimulus load of 10^{-12} g, far lower than the earlier statistical analyses based upon excitatory responses suggests.

3.4.5.2 Further Analysis

Since the ALN shows inhibition, the net spikes s^{-1} method used to quantify the response of excitatory neurons is not particularly suitable since the time over which the net spikes are calculated uses an increase in the firing rate to determine the start (although it is possible to use the net spikes s^{-1} method using a decrease in the firing rate, Appendix B). Where there is a decrease in the firing rate, it is unclear exactly when the inhibition starts, so the start of the time interval over which the net spikes s^{-1} are calculated is unclear. Using the last "background" spike as the start of the inhibition onset is not accurate since the timing of the last spike depends on the "state" of the neuron i.e. if the neuron has recently

produced a spike then the perceived start time of the inhibition may appear later than that of a neuron that is just about to produce a spike (see Figure 3.8).

Another possible method for creating dose response curves is to use the instantaneous spike frequency (ISF). To calculate the ISF, a list of inter-spike intervals are generated and then inverted. The time of each ISF is taken as the mid-point of the time between the two spikes it relates to. Since inhibition corresponds to a decrease in the firing frequency, the dose response curve is created using the minimum ISF in a time period (starting at the stimulus onset and ending 1.5 s later). The dose response for the inhibition ALN, created using this method, is shown in Figure 3.7 b).

Although insufficient examples of such responses precluded detailed t-test or ROC analysis, it must be concluded that the previous estimates of sensitivity enhancement are extremely conservative. Indeed, such inhibitory responses would suggest a sensitivity boost closer to 5 orders of magnitude is possible in this system.

3.4.6 Theoretical Model of Convergence

Because the convergent architecture of the early olfactory pathway in this animal is highly organized topographically, it was possible to compare this empirical boost with a theoretical model of signal convergence, to test the hypothesis that the sensitivity boost results from the convergence of elevated firing rates alone. It can be assumed that ORN spike trains follow an inhomogeneous Poisson point process, with a firing probability that is elevated during stimulus presentation. The case where n such ORN spike trains were integrated over time, at a single site of convergence such as a glomerulus was considered (see Figure 3.9 a)).

The mean firing rate and variance for an ORN both pre- and peri-stimulus are denoted by μ_{pre}, σ^2_{pre} and μ_{peri}, σ^2_{peri} respectively. If an excitatory response to the stimulus is assumed, then the signal-to-noise ratio (SNR) for the response of a single ORN is given by

$$SNR_{ORN} = \frac{1}{2}\frac{(\mu_{peri} - \mu_{pre})^2}{\sigma^2_{peri} + \sigma^2_{pre}} \tag{3.1}$$

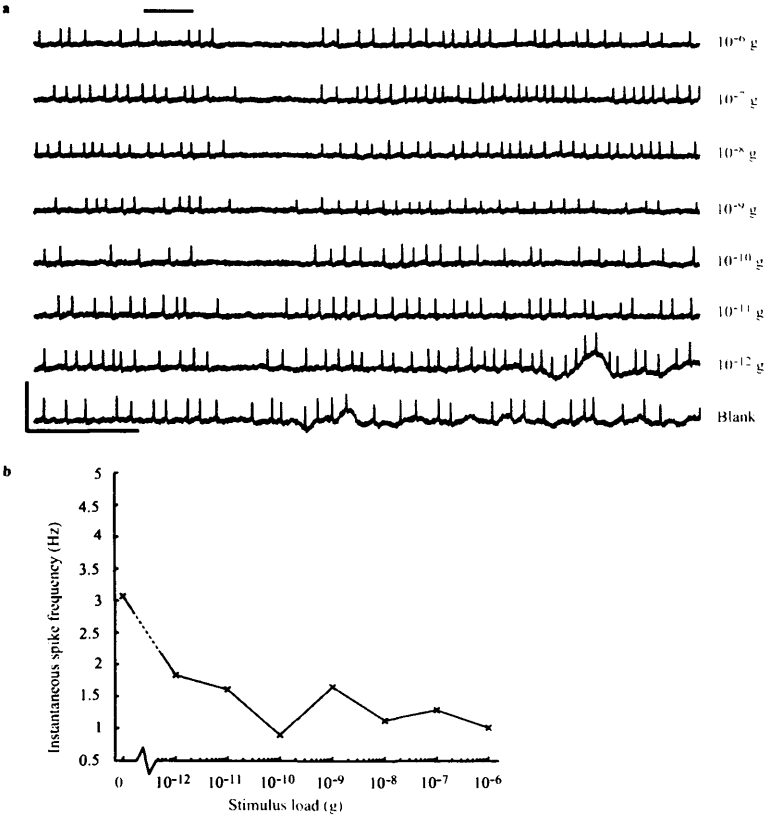

Figure 3.7: Antennal Lobe Neurons Show Inhibitory Response to Lower Stimulus Loads. a) Recordings From an Inhibited ALN. An ALN responding to the major pheromone component, Z9, E11-14:OAc, at different stimulus loads. The neuron shows a response to the stimulus presentation with a suppression of the background firing rate. This inhibition is evident down to a stimulus load of at least 10^{-12} g (horizontal bar shows stimulus duration, horizontal scale bar is 1 s, vertical scale bar is 100 mV). b) Dose Response Curve. To measure the amount of inhibition, the ISF is used (the reciprocal of the time between one spike and the next). The ISF is calculated over a period of time similar to that of the net spikes calculation, and the lowest ISF is used to create the dose response curve. Lower spike frequencies indicate higher levels of inhibition. This neuron shows lower spike frequencies for all stimulus loads tested compared to that of the blank stimulus.

Under the simplifying assumption (which is later abandoned) that the role of a secondary neuron (ALN) is to emit a spike for every spike received at a point of integration where n ORNs converge on a single glomerulus, then the firing rate for the ALN becomes $n\mu_{peri}$ during the stimulus and $n\mu_{pre}$ beforehand. Moreover, since for a Poisson process the variance of the firing rate is equivalent to the mean, then the SNR for the ALN becomes

$$SNR_{ALN} = \frac{n}{2}\frac{(\mu_{peri} - \mu_{pre})^2}{\mu_{peri} + \mu_{pre}} \tag{3.2}$$

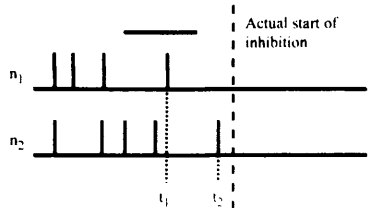

Figure 3.8: Perceived Time of Inhibition Onset. Horizontal bar represents the stimulus duration. Dashed line represents the time at which the true inhibition starts. Neuron n_1 is just about to fire a spike but becomes inhibited. Neurons n_2 fires a spike just before becoming inhibited. Times t_1 and t_2 are the perceived times of the onset of inhibition for the two neurons.

Since the SNR of the ORN does not depend upon n the enhancement in SNR, or SNR_{ALN}/SNR_{ORN} follows n. The ability of an ideal observer to detect the stimulus depends upon the psychophysical d' which is measured in terms of the square root of the SNR (Dayan and Abbott, 2001). Thus the discriminability (sensitivity) enhancement of the system follows $\sqrt{n}$ directly.

Using a simple computational model, it was possible to simulate the convergence of large numbers of ORNs. Each ORN consisted of 1000 time bins, each of 1 ms in duration. A random number generator was used to determine whether each bin contained a spike or not, according to a probability set such that the pre-stimulus firing rate (first 500 time bins) was $^1/_3$ that of the post-stimulus firing rate (last 500 time bins). At the site of convergence, the input ORN spikes were compressed into a single spike train, again of 1000 time bins, containing either a 0 (no ORNs had a spike in this time bin), or a 1 (one or more ORNs had a spike in this time bin). These convergence calculations were performed 500 times to calculate the psychophysical d'. When d' is plotted against $\sqrt{n}$ the relationship between them can be modelled by a linear best fit (see Figure 3.9 b)).

This improvement in discriminability can potentially be translated directly into a sensitivity enhancement for the system. Therefore, to achieve the observed boost of 3 orders of magnitude through pooling of firing rates would require 10^6 ORNs converging onto a single glomerulus, whereas a boost of 10^5 would require 10^{10} ORNs.

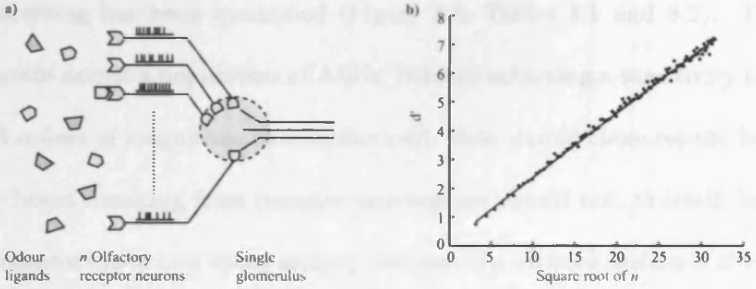

Figure 3.9: Theoretical Convergence Model. a) Spike Train Integration at a Glomerulus. A schematic view of receptor convergence onto a singe glomerulus in the early stages of olfactory processing. Odour ligands bind to receptor sites on ORNs producing above background firing rates. Axons from n ORNs expressing the same receptor gene converge onto a single glomerulus in the MGC. b) Boost in Discriminability. A simple computational model of the theoretical model shows that when the ORNs are created as Poisson point processes, and the output neuron simply sums the input spikes, the boost in discriminability of the system follows $\sqrt{n}$. For each n (up to 1000), 500 trials were performed to calculate d'. Dots show the actual values of d', with boost in sensitivity following $\sqrt{n}$, shown by the line of best fit.

3.5 Discussion

3.5.1 Receptor Neuron Pooling and Sensitivity Enhancement during Early Olfaction

Within the olfactory system, glomeruli act as sites of convergence for large numbers of ORNs (Christensen et al., 1996). This convergence is conserved across many species of both invertebrates and vertebrates (Allison and Warwick, 1949; Bhatnager and Kallen, 1975; Gemne and Døving, 1969), suggesting a fundamental role in the detection of chemical cues. Electrophysiological recordings taken from olfactory bulb (OB) of frog compared with those at the receptor level show increased sensitivity to odours, which can be interpreted as a sensitivity enhancement (Duchamp-Viret et al., 1989), which is also observed in the AL of the male moth *Agrotis segetum* (Carlsson et al., 2002; Hartlieb et al., 1997) and male cockroach *Periplaneta americana* (Boeckh and Ernst, 1987). Taken together, these observations suggest that a key property of massive convergence in the olfactory system is signal amplification, resulting in increased sensitivity to chemical stimuli.

By conducting a signal detection theoretic analysis of firing rates in the olfactory pathway in this chapter, the boost in detectability of the olfactory signal occurring during early

olfactory processing has been quantified (Figure 3.5, Tables 3.1 and 3.2). This boost is shown to operate across a population of ALNs, reliably achieving a sensitivity enhancement in excess of 3 orders of magnitude in stimulus load. How should these results be explained? A sensitivity boost resulting from receptor convergence should not, in itself, be surprising. By pooling receptor input over many sensory neurons, the nervous system is able to decrease uncertainty, which can be translated as an overall lowering in the detection threshold, so that higher centres of the nervous system can robustly interpret chemical cues. Unlike time averaging, such a spatial integration of receptor input incurs no time penalty; in principle this can be performed in the same time required to readout the signal from a single neuron.

3.5.2 Olfactory Sensitivity Boost is Beyond Pooling of Firing Rates

Since glomeruli are innervated by axons projected by ORNs with identical tunings, the question of whether the sensitivity boost observed could derive from a straightforward integration of firing rates at these structures was asked. Theoretical considerations of integration of Poisson spike trains at a single site of convergence, such as at an individual glomerulus, show that discriminability of a stimulus must follow the square root of the number of convergent spike trains (see Section 3.4.6). Thus, to achieve the observed boost of 3 orders of magnitude through pooling of firing rates would require 10^6 ORNs converging on a single glomerulus, whereas a boost of 10^5 would require 10^{10} ORNs. The known anatomy of the sensilla and MGC in the moth, where axons from ORNs expressing the same olfactory receptor converge (Berg et al., 1998; Hansson, 1995; Hansson et al., 1992; Ochieng' et al., 1995; Todd et al., 1995), puts this figure closer to 10^4, suggesting that a maximum of 2 orders of magnitude boost could be attained by the system through pooling of firing rates alone. ALNs showing inhibitory responses in excess of 5 orders of magnitude lower than the detection threshold measured at the periphery (Figure 3.7) were found, in itself suggesting that an extreme level of sensitivity enhancement exists in this system.

Before discussing possible mechanisms that could conceivably account for the boost in detectability of the stimulus, it is first necessary to assess the assumptions made throughout this chapter.

3.5.3 Assessment of Assumptions

Throughout this chapter, a number of assumptions have been made, about both ORNs and ALNs, and their connectivity. This section examines these assumptions, with brief discussions about the relevance of each, and what has been done to try to discount errors associated with the assumptions.

1. All ORNs responding to the stimulus are of the same type and have the same sensitivity i.e. belong to a homogeneous population.

   As mentioned in Section 3.4.2.2, ORNs are likely to express only a single type of olfactory receptor (Buck, 1996; Chess et al., 1994; Ngai et al., 1993), and thereby certain specificity (Carlsson et al., 2002; Hansson et al., 1992; Ochieng' et al., 1995). When stimulated with pheromone, it can be expected that the population will behave in a stereotypical manner. Indeed, investigations have previously been carried out (including Ljungberg et al., 1993), which show that the detection threshold for ORNs sensitive to the major pheromone component were similar to those found in this chapter, and that pheromone-detecting ORNs are highly specialized (Hansson, 1995 and references therein). It can be concluded that the ORNs all show a maximal response to the stimulus and can be perceived as a homogeneous population when calculating the detection threshold for a specific stimulus.

2. ALNs responding to the stimulus do not behave in the same way i.e. belong to a heterogeneous population.

   As mentioned in Section 3.4.2.3, ALNs innervating the same glomerulus in the MGC may display different degrees of specificity. Unlike ORNs, the specificity of ALNs

does not arise from a direct interaction with odour molecules, but rather from the interconnections between neurons within the AL. This property of ALNs has been studied (see Anton and Hansson, 1995) with the conclusion that ALNs may respond to several different stimuli, with a maximal response to only one of them. It should therefore be acceptable to assume heterogeneity in the ALN population, with a subset of maximally responding ALNs used for the detection threshold comparison.

3. The ORNs recorded from pass the signal to the same glomerulus that the ALNs read from.

   The following points indicate that both types of neurons innervated the same glomerulus:

   (a) *Spodoptera littoralis* has only three glomeruli in the MGC, one each receiving axons from ORNs responding to the major and minor components of the pheromone, and one receiving axons from ORNs responding to a behavioural antagonist (Ochieng' et al., 1995);

   (b) Calcium imaging experiments show that ORNs responding to the major pheromone component generally innervate a single MGC glomerulus (Carlsson et al., 2002);

   (c) Innervation sites used to locate ALNs in the MGC that respond to the stimulus were based on glomerular representations (Carlsson et al., 2002; Sadek et al., 2002) and calcium imaging (Carlsson et al., 2002).

   However, without having taken simultaneous recordings from ORNs and ALNs, it is not possible to confirm that the ORNs and ALNs recorded from were innervating the same glomerulus.

4. PNs and LNs can be grouped together as ALNs.

   The grouping of PNs and LNs into ALNs was made due to the lack of morphological evidence available to identify the individual types of neurons. What this assumption really results in is that it is not possible to determine whether the boost seen at the AL is confined to the AL (i.e. is a result of the synaptic connectivity and is only seen in LNs), or whether the boost is transported to other areas of the moth brain (i.e. is carried in the PN output response). Both types of neurons are, however, located in the AL, and are found to innervate glomeruli, including those in the pheromone processing subsystem. It is therefore still possible to demonstrate an enhancement in the sensitivity of the olfactory system between the antennae and the AL.

5. An inhibited ALN can be used to describe further boost than is seen in excitatory ALN responses.

   In *Spodoptera littoralis*, the LNs can have an inhibitory effect on other neurons, but PNs and ORNs are only able to excite other neurons. For an ALN to show an inhibitory response to a stimulus, an LN must have responded to the stimulus and inhibited the ALN via synaptic connections. This is similar to the point above, and is nevertheless a demonstration of the enhancement in sensitivity between the antennae and the AL.

6. The coding strategy of interest for both the ORNs and the ALNs is a rate code.

   For the purposes of this chapter, it has been assumed that the presence of a stimulus is portrayed in the firing rate of a neuron, usually by an increase in the firing rate (with the exception of the inhibited ALN). However, theoretic and experimental investigations suggest that sensory information could be transmitted as a temporal code (Rieke et al., 1997; Theunissen and Miller, 1995). This idea is discussed in Section 3.5.4.

7. The time period over which the response was calculated was of relevance to the biological system.

   This chapter assumes a rate code (see point 6), and to calculate this, an integration period is required. The integration period used is described in Section 3.3.2 and was the time period in which individual neurons showed an elevated firing rate compared with the background firing rate. It was necessary to try to discount the possibility that considering alternative timescales, possibly more relevant to the animal, could enhance the discriminability of the stimulus. Section 3.4.4 describes this procedure, with the result being that alternative integration times do nothing to improve stimulus discriminability. Without knowing exactly what the biological system determines as important, it is not possible to discount the possibility that lower detection thresholds, particularly in the ORN responses, could not be found if the ideal integration period were known.

8. The theoretical model used is a good representation of the biological system.

   As an initial theoretical investigation, ORNs were assumed to be inhomogeneous Poisson point processes. The real ORNs recorded from were found to have coefficients of variation between 0.5-1.5, suggesting non-Poisson properties. The assumption that the output neuron emits a spike for every spike received, allows for a simple theoretical model, but is not a good representation of an ALN. These issues are addressed in more detail in Chapter 5.

These were the main assumptions made whilst assessing the detection thresholds of ORNs and ALNs, and during the investigation of enhancement in stimulus discriminability. After considering these assumptions, it is still possible to realise that the boost in stimulus detectability, between neurons at the antenna and those in the AL, is greater than can be accounted for by simple summation of firing rates of receptor inputs. This suggests that alternative coding and/or readout mechanisms are operating in this system, which require further scrutiny.

3.5.4 Further Hypotheses

3.5.4.1 Further Sensitivity at the Periphery

While still assuming a rate code at the periphery, it is not possible to completely discount the possibility that ORNs more sensitive to the stimulus exist. Although this is unlikely since thousands of sensilla have been recorded from in previous studies (where similar detection thresholds were found), recordings have mainly been taken from a single type of sensilla (the trichoid sensilla). If ORNs more sensitive to the pheromone component used here do exist, there is a possibility they may be housed in other sensilla not yet so rigorously investigated.

When using the current experimental set-up to stimulate the antenna, it is possible that only a small section of the antenna actually receives the puff of air containing the pheromone. At low stimulus loads, not all sensilla may receive enough pheromone molecules to cause a visible response in the ORN housed inside it. In fact, for two of the ORNs used in this chapter, the recordings at lower stimulus loads suggest that these neurons didn't always receive enough of the stimulus to provide a response in each repetition. For these two ORNs, five recordings were taken at each stimulus load (blank, 10^{-9} g, 10^{-8} g, 10^{-7} g, 10^{-6} g and 10^{-5} g) and visual inspection suggests that responses to the stimulus presentation are seen at 10^{-8} g in 1 out of 5 and 2 out of 5 recordings. Examples of recordings both with and without a response are shown in Figure 3.10 a) and b). This hypothesis, that sometimes an ORN will receive enough stimulus molecules to evoke a response, but other times it doesn't, suggests that although the overall population average detection threshold is 10^{-7} g, ORNs may be more sensitive. Although this reduces the difference in threshold between ORNs and ALNs, there would be far fewer responding neurons converging at the glomerulus than assumed in the original investigation of the population, so the boost in discriminability remains intact. This hypothesis is also investigated in Chapter 5.

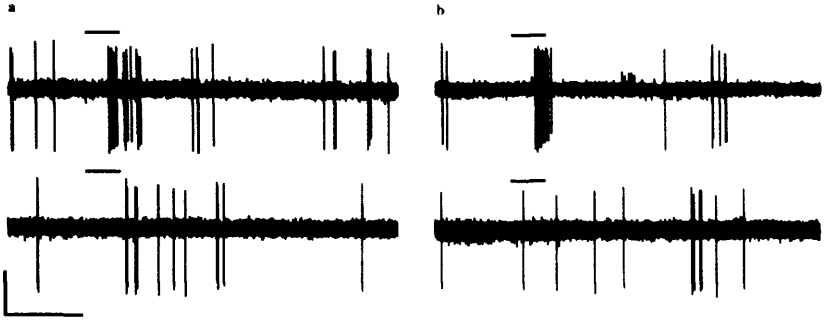

Figure 3.10: Example of Possible 'Lucky' ORNs. a) and b) show two recordings from two different ORNs, all with a stimulus load of 10^{-8} g (stimulus duration represented by short horizontal bars above each recording). Stimulus load tested is below the threshold determined when considering the whole population of ORNs, but the top row shows an increase in firing rate in response to the stimulus, whereas the bottom row does not (although later presentations at higher stimulus loads still showed responses, data not shown). Horizontal bar: 1 second, vertical bar: 1 Volt.

3.5.4.2 Temporal Patterns at the Periphery

Throughout this investigation, only the pooling of elevated firing rates has been used to assess the boost in detection performance. The comparison of a theoretical model of convergence with the empirical data suggests that other coding mechanisms may be in operation to further boost detection performance. If the ALNs were sensitive to temporal information in the form of precisely timed spikes, or patterns of spikes generated at the periphery, these spikes would not necessarily lead to large excursions in ORN firing rates at low stimulus loads. In this coding scheme, an ALN would need to detect a specific, repeatable pattern of spikes received from the ORN population. Although it can be seen that the overall structure of firing rates of ORNs follows a pattern, it has not yet been ascertained whether ORN spike trains over repeated trials, at stimulus loads below the current detection threshold, consist of an obviously repeatable temporal structure. This will be further investigated in Chapter 4.

3.5.4.3 Synchronisation at the Periphery

Another method by which sensitivity enhancement could be further boosted is through individual or clusters of spikes generated by the ORNs, which are synchronised in time during

stimulus presentation. This encoding strategy, if only required to indicate the presence of a stimulus, could be achieved by each ORN transmitting a single spike, synchronised with spikes from a population of ORNs, again having little effect of the overall firing rate of the sensory neurons. The ALN would need to be sensitive to these coincident events over the ORN population innervating its glomerulus, as investigated in Chapter 5.

In order to achieve the synchronous firing of spikes from a population of neurons, a mechanism would be required to produce this synchrony. Two possibilities exist: the dynamics of the stimulus, common to all receptor neurons, is sufficient to time-lock the ORN activity with the degree of precision necessary for synchronous detection; ORNs are able to synchronise their activity through other means. In the case of the moth, chemical stimulus dynamics up to 25 Hz (Justus et al., 2002b) have been measured in behaviourally relevant conditions and this could in principle induce highly synchronised stimulus onset times across the numerous sensilla. In mammals the signal received at different compartments of a single glomerulus has also been shown to be synchronised within the time resolution of two-photon calcium imaging (Wachowiak et al., 2004). It is therefore possible that the common stimulus dynamics in the moth are able to sufficiently synchronise ORN activity to permit downstream detection (further investigated in Chapter 4). Any other method of synchronisation seems unlikely, since there is no evidence that ORNs across distinct sensilla are able to coordinate their activity to achieve the necessary synchronisation.

3.5.4.4 Models of Glomerular Convergence Sites

Within this chapter, the theoretical model of convergence uses Poisson ORN spike trains, with a simple summation of these at the convergence site. The result of this is that the enhancement of stimulus discriminability should follow the square root of the number of converging inputs. When compared with the biological system, the enhancement observed exceeds that predicted by the theoretical model. The theoretical model is, however, very basic, with the possibility that using more realistic firing statistics for the input ORNs,

and a more realistic model of an output neuron could conceivably give rise to a higher enhancement than predicted by the theoretical model.

A more biologically realistic computational model may still not be able to account for the enhancement found in the biological system when considering a passive site of convergence. However, research has shown that glomeruli are highly complex units innervated by several types of neurons (Distler et al., 1998). Innervating neurons form within-glomerulus networks, while also forming neural connections to other glomeruli and to other brain areas. It is, however, most likely in the within-glomerulus circuitry that the signal enhancement mechanisms should be sought. In a modelling study of LNs, one conclusion was that at low concentrations of odours these neurons most likely work in the intra-glomerular space (Christensen et al., 2001). It was also shown how PNs, which constitute the main output from the AL, innervating the same glomerulus go into synchrony at a much higher degree than those present in neighbouring glomeruli. The PNs of neighbouring glomeruli could even inhibit the activity of each other (Lei et al., 2004). Chapter 5 looks at ways in which synaptic connections between neurons within a glomerulus may result in the formation of an active site of convergence, and so enhancing the detection capability of the system.

3.5.5 Temporal Patterns in the Antennal Lobe

For the detection thresholds investigation in this chapter, only a rate-based metric, calculated for individual ALNs, was considered. Often, however, ALN responses to olfactory stimuli are complex, consisiting of excitation and inhibition, and outlast the duration of the stimulus presentation (Laurent and Davidowitz, 1994; Laurent et al., 1996; Lemon and Getz, 2000). Indeed, neural activity within the antennal lobe appears to be both spatially and temporally complex (Gelperin, 1999; Laurent, 1999; Sachse and Galizia, 2002), with individual neurons displaying odour-specific spike patterns (and the same odour evoking different patterns in different neurons; Laurent et al., 1996). Recordings across multiple neurons suggest that specific patterns of synchronous firing across a population of neurons

may encode different aspects of the olfactory stimulus (Christensen and Hildebrand, 2002; Christensen et al., 2000; Hansson and Christensen, 1999; Laurent et al., 2001; Lei et al., 2002). While these ideas are beyond the scope of this thesis, it is conceivable that temporal responses to pheromone stimuli in ALNs may result in lower detection thresholds than found here, further enhancing the detection performance.

3.6 Conclusions

By conducting detailed signal theoretic investigations of discriminability at the first two stages of the olfactory pathway a significant sensitivity boost has been found. Comparing this boost to a theoretical model of signal convergence, this boost is beyond that achievable by summations of elevated firing rates over a pool of receptor neurons. The assumptions made throughout the investigation were discussed, with the conclusion being that none caused the overall outcome to deviate significantly. A variety of hypotheses to explain the boost observed were suggested, from different coding schemes to alternative readout mechanisms, with the aim to investigate some of these further in the following chapters.

Chapter 4

Temporal Encoding at the Periphery

4.1 Chapter Overview

The main aim of this chapter is to perform an initial investigation into the encoding scheme used by the ORNs at the periphery of the olfactory system. Chapter 3 suggests that the boost in signal detection between ORNs at the antennae and the ALNs in the AL is greater than that expected when assuming a rate-based encoding scheme at the ORNs, and a passive site of convergence of these ORNs at the ALNs. One solution that could be employed to explain this boost by the olfactory system is the use of temporal encoding at the periphery (examples of which are described in Sections 3.5.4.2 and 3.5.4.3), producing a code not detectable using a firing rate metric.

Several repeats of electrophysiological recordings from a number of different ORNs were taken by collaborators in Sweden (see Section 2.2) in response to the major pheromone component (at stimulus loads of the rate-based detection threshold and lower), and spike trains were generated as detailed in Section 2.3. Stimulus correlation techniques were used to initially assess the ORN encoding scheme, the results of which could not discount the possibility of a temporal encoding scheme. Patterns of spikes from repeat recordings of individual ORNs were then studied using Unitary Events Analysis (UEA) but most were found to be unreliable. The problem of non-stationarity of spike trains was investigated.

Unitary Events by Moving Window Analysis (UEMWA) was used to investigate patterns of synchronised spikes across a population of ORNs. Overall, the results were inconclusive, mainly due to the non-stationarity of the spike trains, combined with low background firing rates of the current data set causing complications when using these techniques.

Assumptions relating specifically to the methods used throughout the chapter are discussed, with the main assumption, that of stationarity of spike trains, being violated resulting in an unreliable outcome. Further ideas for investigation are suggested, some based on the current data set, and others requiring specific data sets not yet available.

4.2 Introduction

In Chapter 3, the results show that the boost in sensitivity between the first two olfactory processing stages, ORNs at the antennae and ALNs of the AL, is beyond that achievable by the summation of elevated firing rates from the convergence of the population of ORNs. If, however, the ALNs were sensitive to stimulus information encoded temporally by the ORNs, the spikes from this type of encoding scheme would not necessarily lead to large excursions in ORN firing rate, particularly at lower stimulus loads, which would not have been identified as part of the investigation conducted in Chapter 3.

Temporal encoding could be achieved by the ORNs in several ways. Individual ORNs could produce patterns of spikes, differing according to the stimulus load and/or stimulus odour quality. For these patterns to be transferred to the AL, ORN spike patterns would need to be highly repeatable over multiple stimulus presentations, with accurate spike timings to allow for recognition of the pattern. Another scenario requires spikes (individual or in clusters) from ORNs to be synchronised across the population. The ALNs would need to be capable of detecting these coincident events, and the synchronisation would also need to be repeatable over multiple stimulus presentations. The purpose of this chapter is to investigate the possibility of temporal encoding at the ORNs, with Chapter 5 investigating possible mechanisms at the AL.

Assuming then, that these temporal encoding schemes exist in this olfactory system, how might they be identified? A set of criteria have been established for the qualification of an encoding scheme as either a rate-based or a temporally-based code (Theunissen and Miller, 1995), according to the degree of correlation between the spike trains and the stimulus. This method has been applied in the olfactory system of the cockroach, with the finding that the ORNs used a rate-based encoding scheme, which is transformed into a temporal encoding scheme at the PNs (Lemon and Getz, 2000). While this procedure can be used to identify the presence of a temporal encoding scheme, it does not indicate the structure of the encoding scheme.

To further assess temporal encoding schemes, several methods have been developed including joint-PSTH (jPSTH; Aertsen et al., 1989), sliding window cross-correlograms (Laurent and Davidowitz, 1994), cost-based metrics (Victor and Purpura, 1996, 1997) and the detection of stable spike patterns (Fellous et al., 2004). While these methods serve their intended purpose, they only compare one neuron with another, and are not capable of detecting synchrony across a larger population of neurons. A method more suited to this form of analysis is unitary events analysis (Grün, 1996; Grün et al., 2001a,b). UEA can be used to investigate temporal patterns across repeat trials of single or multiple neurons, with the number mainly restricted by the computational time.

To investigate the encoding scheme utilised by ORNs, first, stimulus correlation techniques will be used in this chapter to validate the possibility of the encoding scheme being temporal (i.e. the timing of the spikes play an important role in stimulus detection). UEA is then used in two different ways: firstly, to investigate the reliability of patterns of spikes over repeated stimulus presentations for individual ORNs, and secondly to investigate the degree of spike synchrony across a population of ORNs across repeated trials. The analyses performed in this chapter find that although the ORN encoding scheme could be temporal, the investigation of the specific temporal structure is inconclusive.

4.3 Data Analysis Methods

4.3.1 Electrophysiological Recordings

Collaborators in Sweden performed electrophysiological recordings from many ORNs (see Section 2.2). The stimulus used was the major component of the female sex pheromone, Z9, E11-14:OAc. The stimulus duration was 0.5 seconds, with stimulus loads ranging from 10^{-5} g to 10^{-9} g. Spike timings were extracted from each recording according to the methods described in Chapter 2. For the preliminary analysis (Section 4.4.1), recordings from ORNs were only used if there were at least four repeats (with a maximum of the first five repeats actually being used) for at least two stimulus loads and blank presentations, resulting in the analysis of 15 different ORNs.

4.3.2 Stimulus Correlation with Spike Count

The first part of the preliminary analysis in this chapter is used to assess how the stimulus is encoded (either as a rate or a temporal code; Lemon and Getz, 2000). Recordings from ORNs were organised according to the stimulus load presented, then each spike train was divided into 50 ms time bins (the encoding time window, see Section 4.5.2), and the total number of spikes that occurred during each bin was counted. The average spike frequency of each ORN in response to the blank presentations was subtracted from each response to odour. Adjusting the responses in this manner removes any mechano-sensory response to the change in airflow caused by stimulus presentation. The portion of each response analysed started at the stimulus onset, and continued for 1000 ms (the integration time window, see Section 4.5.2).

The adjusted sequences of spike counts per 50 ms bin from each spike train, in response to two stimulus loads, were then used to calculate the first two principal components (PCs; see Appendix A) using MATLAB's `princomp` function. The PCs describe the binned spike trains according to the degree of variability explained, with the first PC indicating the

variability due to change in rate, and the second PC (and higher) indicating other coding mechanisms are responsible for producing the variability (Lemon and Getz, 2000). The PC scores (or coefficients) describe the projection of the adjusted spike trains onto the new subspace formed by the first two PCs, resulting in two PC scores per spike train. These PC coefficients were plotted against the total number of spikes found in the adjusted (binned) spike trains to assess the degree of correlation (see Figure 4.5 for examples; Lemon and Getz, 2000).

4.3.3 Comparison of Principal Component Coefficients

The second part of the preliminary analysis in this chapter is again used to assess the coding strategy at the periphery. In this case, recordings were analysed for all stimulus loads for individual ORNs. The same numbers of spike trains for each stimulus load were first split into 50 ms time bins and adjusted by subtracting the average number of spike in each window in response to the blank presentations (as for Section 4.3.2). Five of these adjusted spike trains were sampled uniformly (with replacement) from the pool of all the recordings from an individual ORN. The first two PC coefficients were found, and the sampling procedure was repeated 1000 times. This resulted in two distributions, one for each of the two first PC coefficients, with each containing 5000 coefficients. The same procedure was then carried out for the stimulus loads individually, resulting in two distributions for each of the stimulus loads presented to the ORN. These distributions were generated for each of the ORNs.

4.3.4 Unitary Events Analysis

The preliminary analyses described in Sections 4.3.2 and 4.3.3 are used to assess the possibility of a temporal encoding scheme at the periphery, but give no information on the structure of the coding strategy. UEA was developed by Grün et al. (Grün, 1996; Grün et al., 2001a) to search for these potential temporal patterns across multiple neurons. This method uses multiple spike trains, converted to binary sequences, and assesses the probability of occurrence of sequences of spikes (within the same time bin but across multiple

spike trains), known as coincidence patterns. If a coincidence pattern has a probability of occurring above a threshold (or significance level), it occurs more often than by chance, and is called a unitary event (UE). To search for temporal patterns at the periphery, the UEA method can be used to find quasi-coincident spikes (spikes that occur at the same time in relation to the stimulus across several repeat recordings to the same ORN).

To conduct UEA, individual spike trains are divided into time bins, and the spikes placed in these bins. Ideally, the length of time bin should be such that each bin contains 0 or 1 spike (bins with more than a single spike are truncated). The result is a binary matrix with rows consisting of binned spike trains, and columns showing binary patterns of spikes across all recordings (see Figure 4.1). Using the initial assumption that each recording is independent (the null hypothesis, Equation 4.1), and the assumption that the spike train firing probabilities are stationary over the duration of the recording, the predicted number of occurrences (n_k^{pred}) of each binary pattern ($\mathbf{v}^k$) can be calculated. The probability, p_i of a spike occurring in any time bin of a recording from neuron i is calculated from an estimate of the firing rate, $p_i = c_i/T$, where c_i is the number of spikes in an interval containing T time steps for neuron i. The probability of occurrence of a binary pattern, $\mathbf{v}^k$, can then be calculated as the product of the probabilities of the individual spiking events:

$$H_0 : P_k = \prod_{i=1}^{N} P\left(v_i^k\right), \qquad \text{with } P\left(v_i^k\right) = \begin{cases} P\left(v_i = 1\right), & \text{if } v_i^k = 1 \\ 1 - P\left(v_i = 1\right), & \text{if } v_i^k = 0 \end{cases} \tag{4.1}$$

The predicted number of occurrences is then the probability of occurrence multiplied by the number of time bins, $n_k^{pred} = P_k T$. The empirical number of each binary pattern, n_k^{emp} can be counted, with patterns of low complexity (containing only a single spike or less) discarded since they don't represent a possible coincidence with spikes from other recordings.

For each empirical number of occurrences of a binary pattern, it is possible to compute the statistical significance of the deviation from independence (the predicted number of

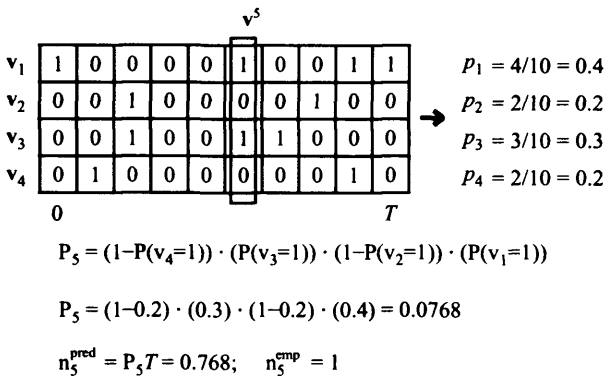

Figure 4.1: Unitary Events Analysis Method. In this case, spikes from four different recordings, $\mathbf{v}_1$ to $\mathbf{v}_4$, have been split into $T = 10$ time bins such that each time bin contains 0 or 1 spikes. Each column now shows a binary pattern of spikes, $\mathbf{v}^k$, across all recordings. The probability of a spike occurring in each time bin for a single recording i can be calculated as $p_i = c_i/T$, where c_i is the number of spike events in T time bins. From these firing probabilities, the probability of occurrence each individual binary pattern, $\mathbf{v}^k$, can be calculated using Equation 4.1. In the example shown, $k = 5$ (from the decimal representation of the binary pattern 0 1 0 1). The number of predicted occurrences of $\mathbf{v}^5$ can be calculated as $n_5^{pred} = P_5 T$, and the actual number of occurrences, n_5^{emp} can be counted from the real data. Figure adapted from Grün et al., 2001a.

occurrences for that pattern). The probability of finding each binary pattern $\mathbf{v}^k$ exactly n_k times can be calculated from the Poisson distribution,

$$\psi(n_k; P_k; T) = \frac{(P_k T)^{n_k}}{n_k!} \exp(-P_k T) \tag{4.2}$$

A logarithmic scaling, $S(\Psi) = \log \frac{1-\Psi}{\Psi}$, of the cumulative probability (Ψ) of finding the observed number of occurrences n_k^{emp} (or more), is used in conjunction with a significance level, S_α, to determine whether this binary pattern occurred more often than expected by chance. If $S\left(n_k^{emp} \,\middle|\, n_k^{pred}\right) \geq S_\alpha$, the binary coincidence pattern is denoted a UE. Values of the threshold, S_α are chosen to correspond to a significance level, α, between 0.05 and 0.01, in order to achieve maximum sensitivity while maintaining a minimum level of false positives. For visualisation, UEs can be plotted as squares on raster plots (see Figures 4.7, 4.8 and 4.9 for examples).

4.3.5 Synthetic Spike Train Generation

The UEA procedure is fairly complex so to ensure correct implementation, a simulation performed by Grün et al.(2001a) is recreated here. Synthetic spike trains were generated as independent homogeneous Poisson point processes with firing probabilities that remained constant for the duration of the spike train. In this case, the probability of observing $(N = X)$ spikes within a time period δt, is governed by the Poisson distribution:

$$P_i(N = X) = \frac{\lambda_i^X}{X!} e^{\lambda_i} \tag{4.3}$$

where λ_i is the mean firing rate for neuron i. Six of these synthetic neurons were created with the following firing rates: neuron 1: 10 spikes s^{-1}; neuron 2: 20 spikes s^{-1}; neuron 3: 15 spikes s^{-1}; neuron 4: 30 spikes s^{-1}; neuron 5: 25 spikes s^{-1}; neuron 6: 15 spikes s^{-1} (as used in Grün et al., 2001a). Spikes had a time resolution of 1 ms, with spike train duration of 100 seconds.

To generate dependent spike trains, the six neurons already created were copied and coincident events were introduced into pairs of these neurons. Coincident events occurred at a firing rate of 1 spike s^{-1} and were created by generating two further independent homogeneous Poisson point processes (using Equation 4.3) to produce coincidences that were randomly distributed in time. The first of these processes was inserted into neurons 1 and 3, and the second into neurons 2 and 5. Neurons 4 and 6 remained independent of each other and the pairs of dependent neurons (as per Grün et al., 2001a; see Figure 4.2).

4.3.6 Test of the Stationarity of Spike Trains

The UEA procedure uses an estimate of the firing rate of each recording to predict the number of occurrences of each binary pattern (Section 4.3.4). It is assumed that the firing rate remains constant (i.e. the rate is stationary) for the duration of the recording. If the recordings are non-stationary, n_k^{pred} will be erroneous, and the comparison with n_k^{emp} will over- or under-estimate the occurrences of UEs rendering the results inaccurate. For

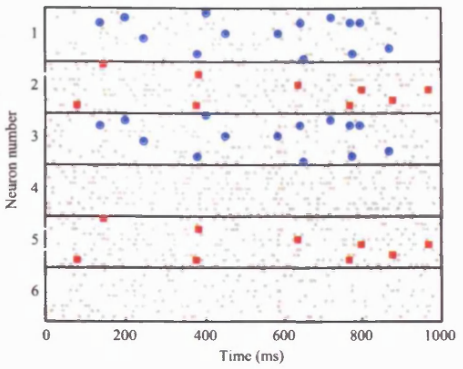

Figure 4.2: Dependent Synthetic Neurons. Six example synthetic neurons were created as independent homogeneous Poisson point processes according to Equation 4.3. The mean firing rates for these neurons were 10, 20, 15, 30, 25, and 15 spikes s^{-1} for neurons 1, 2, 3, 4, 5, and 6 respectively and each neuron consisted of 10 repeats, each lasting for 1000 ms. These independent spikes are shown as grey dots. To introduce dependencies between pairs of neurons, two further independent homogeneous Poisson point process were generated with a firing rate of 1 spike s^{-1}. The first of these extra spike trains was added to neurons 1 and 3 (indicated by blue dots surrounded by a circle) and the other was added to neurons 2 and 5 (indicated by red dots surrounded by a square).

a spike train to be stationary, at the very least, the firing rate should remain constant throughout the duration of the recording. On initial inspection of the ORN spike trains, the presentation of higher stimulus loads (particularly those above the detection threshold) evoked a visible increase in the firing rate indicative of a non-stationary rate around the stimulus onset, whilst in lower stimulus loads (below the detection threshold) this increase was not as evident (and so could be stationary). To further test the stationarity of the lower stimulus loads, the following procedure has been here developed:

1. The spike trains were first converted into either 5 ms or 10 ms time bin representation as per the UEA method.

2. The moving average of each individual representation of a spike train was generated using the difference equation:

$$y_i = \frac{\left(x_i + x_{i-1} + x_{i-2} + \ldots + x_{i-(N-1)}\right)}{N} \tag{4.4}$$

where i is the current time bin and N is the number of time bins required to cover 1 second of the recording (i.e. the window size)

3. The mean (μ_{MA}) and standard deviation (σ_{MA}) of the moving average sequence for the recording section representing the background-firing rate of the neuron were calculated. This section was assumed to be stationary and consisted of the time after which the start up transient of the filter ended (taken to be the same length as the filter i.e. 1 second for this analysis) to the start of the stimulus presentation ($\sim$1 second i.e. size of the window; see Figure 4.3).

4. Spike trains with any section of the moving average sequence (after the start up transient) outside of the $\mu_{MA} \pm (3 \times \sigma_{MA})$ upper and lower bounds were classed as non-stationary. These bounds are not meant as strict guidelines, but used only to imply possible non-stationarity.

Using this method, the stationarity of each spike train can be assessed and sorted into two groups, stationary or non-stationary, accordingly. Spike trains in response to stimulus loads above the detection threshold often exhibit a large increase in firing rate in response to the stimulus presentation, and are found to be non-stationary (Figure 4.3 a)). Neurons appearing to show no response to the stimulus presentation may still be found to be non-stationary (below detection threshold stimulus load, Figure 4.3 b)), but some are found to be stationary (blank stimulus presentation, Figure 4.3 c)).

The window size of the filter used to calculate the moving average sequence can be altered. Window sizes equivalent to 0.25 seconds, 0.5 seconds and 1 second were used and compared, with little difference in the outcome (see Appendix C.2). The number of standard deviations used to determine the upper and lower bound were also varied. Using 2 standard deviations resulted in a few more spike trains found to be non-stationary, while 4 standard deviations resulted in a few more spike trains found to be stationary (see Appendix C.3). The choice of number of standard deviations is not vital since the purpose of this stationarity test was only to imply that not all the analysed spike trains were stationary.

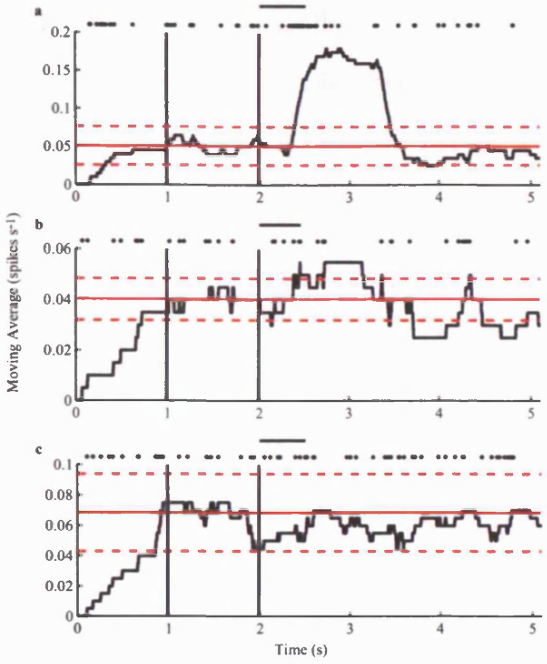

Figure 4.3: Test of the Stationarity of Spike Trains. Black horizontal bar indicates the stimulus presentations. Black dots represent the spike times (in time bins of 1 ms duration). Black vertical lines indicate the section of moving average sequence over which the pre-stimulus mean and standard deviation are calculated. Red horizontal lines show μ_{MA}, the mean pre-stimulus rate (solid line), and the upper and lower limits, $\mu_{MA} \pm (3 \times \sigma_{MA})$ (dashed lines). a) Highly Non-Stationary Spike Train. An example spike train from an ORN in response to a stimulus load of 10^{-6} g (i.e. above the detection threshold). The stimulus presentation elicits a firing rate much larger than the background or pre-stimulus firing rate. This post-stimulus firing rate exceeds the upper limit of $\mu_{MA} + (3 \times \sigma_{MA})$, indicating a non-stationary process. b) Non-Stationary Spike Train. An example spike train from an ORN in response to a stimulus load of 10^{-9} g (i.e. below the detection threshold). The stimulus presentation results in only a slight increase in firing rate. This increase is, however, still sufficient to exceed the upper limit, indicating a non-stationary process. c) Stationary Spike Train. The presentation of a blank stimulus results in an ORN spike train with negligible change in firing rate. The firing rate remains within both the upper and lower bounds throughout, and is consequently regarded as a stationary process.

4.3.7 Unitary Events by Moving Window Analysis

When spike trains are non-stationary, the UEA procedure is not viable since a vital calculation, the prediction of the number of occurrences of a binary pattern, n_k^{pred} (Section 4.3.4), assumes the firing rate remains constant for the duration of the recording. To overcome this problem, when the stationarity of recordings cannot be guaranteed for the duration of a recording, Grün et al. (2001b) propose an extension of the UEA method, namely Uni-

tary Events by Moving Window Analysis (UEMWA). The UEMWA method, described in more detail below, uses a sliding window, over the duration of which the firing rate can be assumed to be stationary. UEA is performed on individual windowed sections of data, allowing each overlapping window to contain a slightly different firing rate.

The main problem with performing UEA on shorter windows of the original data is that the firing rate of neurons is such that too few spikes occur in a single window to allow for the statistical comparison of predicted and observed coincidences. Grün et al. (2001b) provide a solution to this problem by performing the analysis on multiple repeats of the same experiment. Recordings are first divided into time bins as described in Section 4.3.4 and then arranged in trials containing one recording from each neuron (Figure 4.4 a)). The data inside the current position of the sliding window, across all trials, is then concatenated horizontally to produce a 'pseudo-spike train', or section (Figure 4.4 b)), upon which UEA is performed. The window is then slid along the data, usually by a single time bin, where a new section is created on which UEA can be performed. A spike pattern needs to be found to be significant in only a single section for it to be classed as a UE, with the spikes within the UEs again plotted as squares on a raster plot for visualisation.

4.4 Results

4.4.1 Preliminary Analysis of Peripheral Encoding Scheme

4.4.1.1 Correlation of Spike Count with Principal Component Coefficients

To perform this analysis, the initial step was to perform PCA on the binned data from ORN spike trains, as described in Section 4.3.2. The PCs themselves describe a linear combination of the spike trains such that all are orthogonal (see Appendix A). The first two PCs account for more variation than any other pair of PCs, and can act as a two-dimensional subspace onto which the original spike trains can be projected. This projection results in two coefficients, one associated with each of the two PCs, which can be used to assess the encoding scheme according to the above criteria. When considering the first two

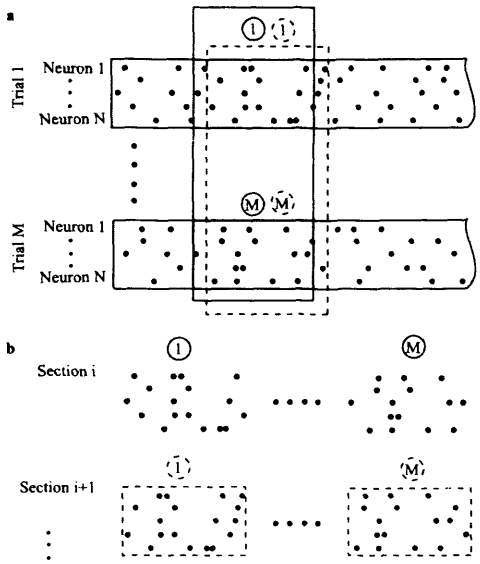

Figure 4.4: Schematic Representation of Concatenation of Windows for Unitary Events by Moving Window Analysis. a) Arrangement of Neurons in Trials. A single recording from each neuron make up each of the M trials. The position of each neuron is preserved across the trials. A time window is used to define a segment of data from each trial. b) Concatenation of Windows. The data within the time windows in each trial are concatenated horizontally producing a 'pseudo-spike train', or section, on which UEA can be performed. The window is then moved to the next time point and the process is repeated to the end of the data set. Figure adapted from Grün et al., 2001b.

PC coefficients of the spike trains from all the ORNs, taken in pairs of stimulus loads, the first PC coefficients are correlated with the spike counts ($r = 0.6825$, $p < 0.0001$, Figure 4.5 a) top), but the second PC coefficients are not ($r = -0.0167$, $p > 0.5$, Figure 4.5 b) bottom).

Spike trains with lower spike counts create a cluster of points with little range in the first PC coefficients, suggesting little contribution to the overall correlation (see Figure 4.5 a) top). To further investigate the degree of correlation of spike count with PC coefficients, the recordings were split into individual stimulus loads, re-plotted, and the correlation re-calculated for each subset. Using the detection threshold found in Chapter 3 (stimulus loads of 10^{-7} g and above are detectable by ORNs if a rate encoding scheme is assumed), the stimulus loads can be split into two groups, those above the threshold (Figure 4.5 b), c) and d)) and those below the threshold (Figure 4.5 e) and f)). Above threshold stimulus

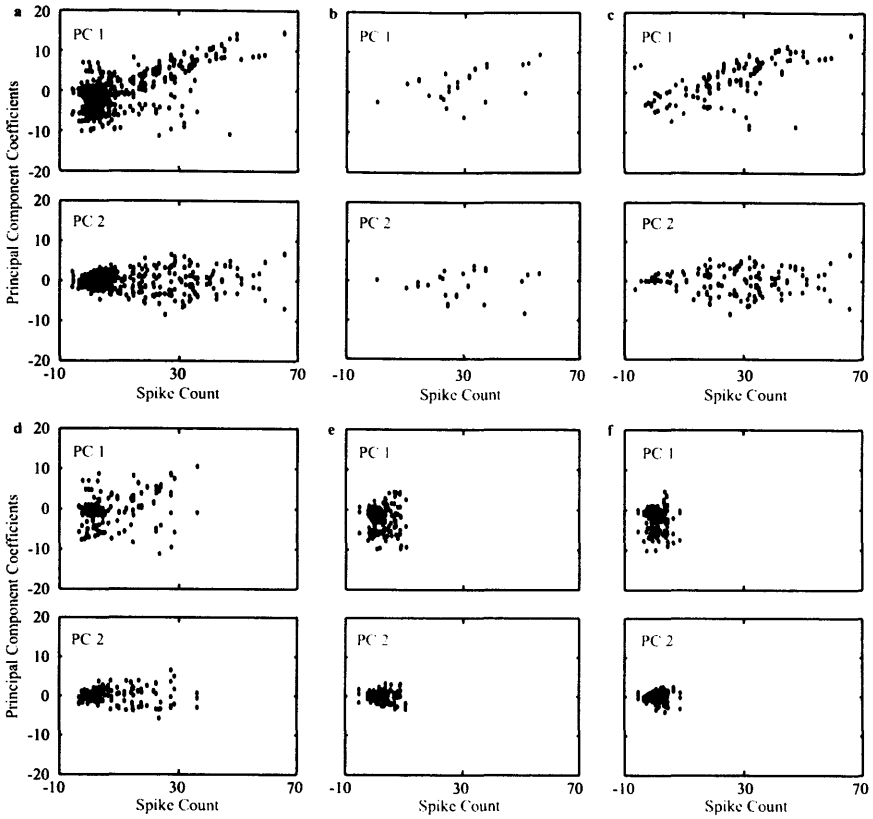

Figure 4.5: Correlation of Spike Count with the First Two Principal Component Coefficients. a) All Stimulus Loads Combined. The coefficients of the first PC, when plotted against the spike counts from individual recordings show a high degree of correlation ($r = 0.6825$, $p < 0.0001$), but the second PC coefficients show no correlation ($r = -0.0167$, $p > 0.5$). The cluster of points with low spike count provides little of the correlation seen with the first PC coefficients, and this is demonstrated in b) to f) where the recordings have been split into individual stimulus loads. b), c) and d) Stimulus Loads Above Detection Threshold. Subplots show stimulus loads of 10^{-5} g, 10^{-6} g and 10^{-7} g respectively. At each stimulus load, different neurons show a range of spike counts, which are correlated with the first PC coefficients, with the degree of correlation decreasing with stimulus load (see Table 4.1). There is no correlation between the second PC coefficients and the spike count for these stimulus loads. e) and f) Stimulus Loads Below Detection Threshold. Subplots show stimulus loads of 10^{-8} g and 10^{-9} g respectively. In both cases, the spike count is not correlated with either of the first two PC coefficients.

loads show significant correlation between the spike count and the first PC coefficients but no correlation between spike count and the second PC coefficients (Table 4.1 a) and b)). Below threshold stimulus loads show no correlation between the spike count and either of the first two PC coefficients suggesting that the individual spike trains at these stimulus loads provide no information on the stimulus presentation.

a) PC 1	Spike Count Correlation with PC 1 Coefficient per Stimulus Load					
	Combined	10^{-5} g	10^{-6} g	10^{-7} g	10^{-8} g	10^{-9} g
Pearson's r	0.68250	0.50641	0.69665	0.36933	-0.04670	-0.04834
p-value	0*	0.00979*	0*	0*	0.53438	0.51812

b) PC 2	Spike Count Correlation with PC 2 Coefficient per Stimulus Load					
	Combined	10^{-5} g	10^{-6} g	10^{-7} g	10^{-8} g	10^{-9} g
Pearson's r	-0.01669	0.05164	0.05201	-0.00403	-0.04829	0.06458
p-value	0.66300	0.80634	0.49554	0.95592	0.52094	0.38773

Table 4.1: Correlation of Spike Count with the First Two Principal Component Coefficients Varies Across Stimulus Loads. Asterisks show correlations that are statistically different from no correlation ($p < 0.01$). a) Spike Count Correlation With First Principal Component Coefficients. When all stimulus loads are combined, the spike count is correlated with the first PC coefficient, and is statistically different to no correlation. When the stimulus loads are considered individually, there is a statistically significant degree of correlation for stimulus loads above the detection threshold (10^{-5} g, 10^{-6} g and 10^{-7} g as found in Chapter 3). For the two stimulus loads below the detection threshold (10^{-8} g and 10^{-9} g), the first PC coefficient is not correlated with the spike count suggesting that the spike train timings provide no information on the stimulus presented. b) Spike Count Correlation With Second Principal Component Coefficients. The spike count is not correlated with the second PC coefficients for any of the individual stimulus loads, or the combined case.

4.4.1.2 Significant Differences of PC Coefficient Distributions from Random Suggests a Temporal Encoding Scheme

Theunissen and Miller (1995; Lemon and Getz, 2000) suggest criteria to assess the encoding scheme employed in neuronal responses. These criteria require the principal components of a measure of neuronal response and state that:

1. If only the coefficient relating to the first principal component is correlated with a feature of the stimulus, the encoding scheme is likely to be rate based.

2. Correlations of a stimulus feature with coefficients of the second principal component or higher indicate a temporally based scheme.

To perform this analysis, the coefficients of the first two PCs were compared to determine if they differed between stimulus loads (here used to represent a feature of the stimulus). Using the sampling method described in Section 4.3.3, distributions of PC coefficients were

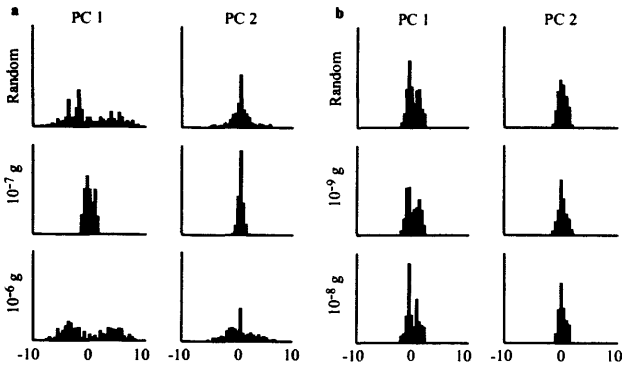

Figure 4.6: Distributions of Principal Component Coefficients. The distributions of pooled coefficients were derived from a sample chosen randomly from the ORN responses with no prior stimulus load information. The distributions of coefficients from individual stimulus loads were generated in a similar manner, but using only ORN responses to the required stimulus load. a) Distributions for Above Threshold Stimulus Loads. Example distributions of PC coefficients generated from the responses of a single ORN to stimulus loads above the detection threshold. The PC coefficients for both stimulus loads are significantly different from the random PC coefficients (Kolmogorov-Smirnov, $p < 0.01$ for both comparisons). b) Distributions for Below Threshold Stimulus Loads. Example distributions of PC coefficients generated from the responses of a single ORN to stimulus loads below the detection threshold. The PC coefficients for both stimulus loads are significantly different from the random PC coefficients (Kolmogorov-Smirnov, $p < 0.01$ for both comparisons).

created for individual stimulus loads and for the pooled mixture for each ORN (see Figure 4.6 for two examples).

Kolmogorov-Smirnov tests were performed to determine whether the PC coefficient distributions for a single stimulus load were statistically different to the distribution for the pooled mixture (Sokal and Rohlf, 1995) i.e. random. If significant differences between the individual stimulus load distributions and the mixed distribution are only found for the first PC coefficients, this suggests that the neurons employ a rate encoding scheme. Significant differences between the individual stimulus load distributions and the mixed distribution for the second PC coefficients suggests a temporal encoding scheme. For all neurons and stimulus loads tested, the individual stimulus load distributions for both PCs were significantly different from the mixed distributions (Kolmogorov-Smirnov, $p < 0.01$ for all comparisons), suggesting the possibility of a temporal encoding scheme.

4.4.1.3 Preliminary Analysis Prompts Further Investigation

The preliminary analysis of the encoding scheme at the periphery was performed to determine whether or not it was possible for the ORNs to encode the presence of a stimulus in a temporal scheme, and not a rate scheme as assumed in Chapter 3. The first test is used to demonstrate that the first PC coefficient is correlated with the spike count, but the second PC coefficient is not. The second test was based on the notion that the first principal component provides information on the variation between spike trains due to rate induced differences, while the second, and higher, principal components provide information on the temporal properties of spike trains. Results indicate that the distributions of both the first and second PC coefficients are significantly different to randomly generated distributions, suggesting a temporal encoding scheme. The overall result is inconclusive in terms of determining whether the type of encoding scheme employed by the ORNs is indeed a temporal scheme, but suggests that further analysis using other methods would be beneficial.

4.4.2 Initial Test of Implementation of Procedure for the Detection of Unitary Events

Although the UEA procedure has been explained in detail (Grün et al., 2001a), with a brief methodology described in Section 4.3.4, this section is used to recreate results from two simulations (Grün et al., 2001a) to ensure the method has been correctly implemented. The first simulation involves six independent model neurons, with the second simulation using the same spike trains, but including the injection of a small number of coincident spikes in some pairs of neurons to indicate dependencies between them.

Spike trains from six simulated neurons were created (modelled as independent homogeneous Poisson point processes, see Section 4.3.5) with differing firing rates (neuron 1: 10 spikes s^{-1}; neuron 2: 20 spikes s^{-1}; neuron 3: 15 spikes s^{-1}; neuron 4: 30 spikes s^{-1}; neuron 5: 25 spikes s^{-1}; neuron 6: 15 spikes s^{-1}). In the first instance, the spike trains from the different neurons were independent and stationary, and lasted for 100 seconds. The UEA

procedure was performed on this group of six neurons using a time bin size of 1 ms, and a p-value of 0.05. The resultant spikes are displayed as raster plots (with the spikes from each neuron plotted as 1 second long sections per row), with any UEs displayed as red squares on these raster plots (Figure 4.7 a) and b) respectively). This independent set of neurons shows three different spike patterns as UEs. Each of these patterns only occurs once, with two patterns containing three spikes each, and the other containing four spikes. The total number of possible patterns containing more than one spike can be calculated as:

$$2^N - \binom{N}{0} - \binom{N}{1} \tag{4.5}$$

For the case of six neurons, Equation 4.5 results in $64 - 1 - 6 = 57$ patterns, and when tested at the $\alpha = 0.05$ significance level it is expected that 2.85 patterns will be marked as significant, closely matching the actual number of significant patterns. Due to the complexity of these patterns, they are likely to be false positives, with the significance found being due to statistical fluctuations (Grün et al., 2001a; Roy et al., 2000).

Dependencies were then introduced into the spike trains by first copying the same six spike trains already created, and adding pairs of simultaneous spikes into neuron pairs 1, 3 and 2, 5 respectively. These coincidences occurred at a rate of 1 spike s^{-1}, much lower than the baseline firing rates, and were randomly distributed in time. The raster plots of both data sets are very similar as expected (compare Figure 4.7 a) and c)). The UEA procedure was then performed on this group of dependent neurons, and again displayed as raster plots. There are now many more UEs than in the independent case, with many corresponding to the injected coincidences. A number of UEs also involve the two independent neurons, but these also included the injected dependent spikes as sub-patterns (Figure 4.7 d)).

The preliminary test of the implementation of the UEA method finds that with six independent neurons, very few UEs are found, and these are likely to be false positives. When dependencies between two pairs of neurons were introduced, far more UEs were

found, with each containing one of the pairs of dependent neurons as a sub-pattern within the UE pattern. Both these findings are comparable with simulations performed by Grün et al. (2001a, Figure 3) suggesting the implementation performs as expected, and can be used on real neuronal data.

4.4.3 Investigation of Temporal Patterns over Repeated Trials in Individual Neurons

In order to determine whether individual ORNs use a temporal encoding scheme to convey the presence of a stimulus, the UEA method described in Section 4.3.4 was performed to detect precisely timed patterns of spikes. UEA has been traditionally described as a method for detecting coincident spikes within a group of simultaneously recorded neurons. For the purposes of this chapter, the UEA method has been applied to multiple spike trains recorded sequentially from a single ORN. In the case of sequential recordings, it is important to ensure the recordings are aligned with each other as accurately as possible, preferably in relation to the stimulus arrival at the neurons. Unfortunately, it is not possible to determine this exact time, so the recordings are actually aligned with the stimulus onset. This could introduce small errors relating to the spike times, but these should be accounted for by adjusting the duration of the time bins (see next paragraph). For this particular analysis, patterns of spikes are expected to occur whenever the same stimulus is presented, although in reality this is not necessarily the case (see Fellous et al., 2004, where different, but stable, spike patterns have been found in response to repeats of a single stimulus). For this analysis, only the first five recordings from each stimulus load for each ORN were used to minimise the possible effects of neuronal adaptation to the stimulus. In addition, recordings to stimulus loads above the detection threshold of 10^{-7} g were not analysed since the presence of a stimulus at these higher loads is adequately conveyed to the AL in the form of a rate code (Chapter 3).

The UEA procedure requires the spikes times to be placed into time bins. The choice of duration of these time bins requires a little consideration. Firstly, the individual ORN

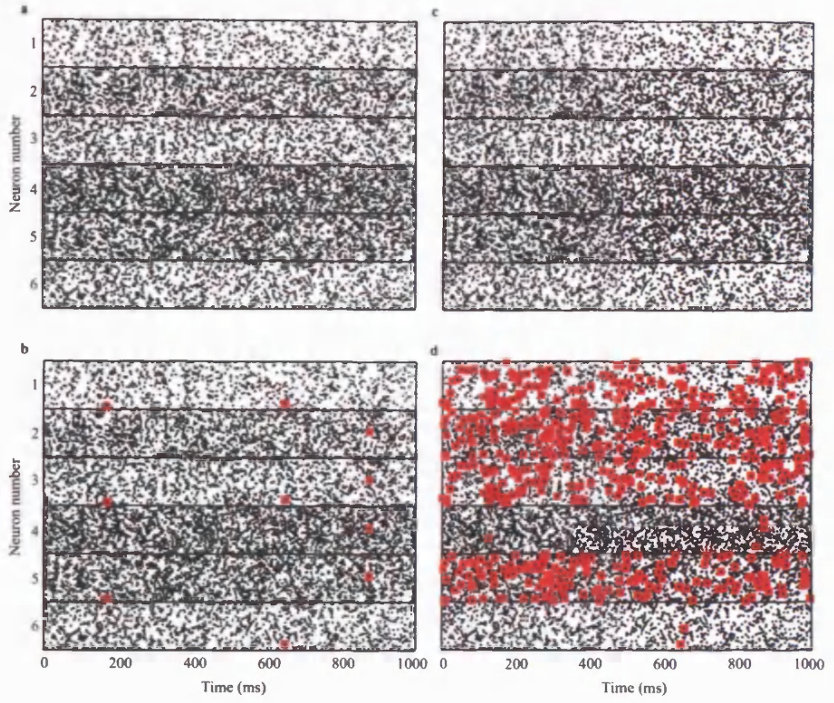

Figure 4.7: Initial Test of Procedure for the Detection of Unitary Events. Synthetic spike trains from six simulated neurons were generated as independent homogeneous Poisson processes. Each spike train lasted for 100 s with the firing rate remaining stationary throughout. The six neurons were simulated such that each had different firing rates: neuron 1: 10 spikes s^{-1}; neuron 2: 20 spikes s^{-1}; neuron 3: 15 spikes s^{-1}; neuron 4: 30 spikes $s^{-1}1$; neuron 5: spikes s^{-1}; neuron 6: 15 spikes s^{-1}. a) Spikes from Six Independent Neurons. Raster plot with each dot representing the time of a spike. Neurons were generated as a 100 s spike trains, but have been displayed as 1 s segments. b) Unitary Events from Six Independent Neurons. Raster plot again shows the spike times, with the addition of the spikes within unitary events displayed as red squares. There are three different spike patterns that occur more times than expected given the firing rates of the neurons. Each of these patterns occurs only once, and contains three or four spikes (in two and one patterns respectively). These patterns are likely to be false positives and due to statistical fluctuations. c) Spikes from Dependent Neurons. The independent spike trains were copied, and simultaneous spikes were inserted into two pairs of neurons, 1 and 3, and 2 and 5. These additional spikes were injected at a rate of 1 spikes s^{-1}, randomly distributed in time, giving rise to a raster plot that on first inspection appears the same as that in a). d) Unitary Events From Dependent Neurons. Again, dots represent the spike times and red squares represent those spikes from unitary events. There are now many more unitary events due to the inserted coincident spikes. Unitary events involving the two remaining independent neurons (4 and 6) also contain either of the pairs of dependent neurons as sub-patterns.

spikes themselves have a duration of ~40 samples, at a sampling rate of 13888 Hz, i.e. ~3 ms (Figure 2.2). Resolution higher than this would not produce further benefit. Neurons are usually considered to exhibit uncertainty, or "noise" in the timing of spikes, although

the amount is not known conclusively (e.g. Diesmann et al., 1999; Mainen and Sejnowski, 1995; Shadlen and Newsome, 1998). Experimental evidence from cross-correlation, jPSTHs and from spike pattern analysis suggests that the timing accuracy of spiking events can be as precise as 1-5 ms (Abeles et al., 1993; Riehle et al., 1997). The possibility of temporal jitter suggests that noisy coincidences occurring across ORN recordings would require larger time bins for detection. Larger time bins, however, increase the possibility of containing multiple spikes, which are then truncated, and possible information is lost. Taking these into account, two time bin widths have been used in this section: an initial bin width of 5 ms (covering the resolution suggested by the spike duration), with a bin width of 10 ms used to encompass patterns with more temporal jitter (initial bin width plus a further 5 ms to encompass precision suggested by Abeles et al. (1993) and Riehle et al. (1997)).

Since only the first five recordings from each ORN and stimulus load were analysed, the total number of patterns of coincident spikes can be calculated from Equation 4.5 as 26 patterns. At a significance level of $\alpha = 0.05$, it can be expected that 1.3 patterns could be found to be significant when they are in fact false positives. To reduce the effect of these false positives, a set of criteria was defined to aid in determining whether the UEs found in the ORN recordings represented true temporal patterns of spikes or not:

1. In order to convey the presence of a stimulus, the stimulus must first occur. Any UEs indicating temporal encoding of the stimulus must therefore occur after the stimulus onset. Any patterns occurring before this time were not considered as patterns of interest.

2. Individual neurons in the AL have been shown to respond within 300 ms of the stimulus onset (Christensen and Hildebrand, 1987; Hartlieb et al., 1997; Sadek et al., 2002). For this to take place, the signal must have been transmitted from the ORNs within this time period. Allowing for a little more flexibility, any UEs occurring later than the stimulus onset + 1 s were assumed to not be involved in triggering these types of behavioural responses and were not considered as patterns of interest.

3. Since the moth does not know in advance when the stimulus onset will be, a single occurrence of a single UE within a single ORN is not sufficient to indicate the presence of the stimulus. In this case, in response to the stimulus, a single ORN produces a single spike (which may occur in repeated trials). This single spike is not distinguishable from the neurons own background spikes, and these events were consequently not considered as patterns of interest. While there is the possibility that these single spikes are synchronised in time with those of other ORNs, this is not considered here.

4. Reliable transmission of the presence of a chemical stimulus suggests that the message should be conveyed by the ORN in response to every stimulus presentation. While this may be unlikely to happen in reality, a limit is set here such that ORNs with fewer than $N/2$ recordings (where N is the total number of recordings from an ORN to that stimulus load) involved in at least one UE (although the UEs themselves do not need to have this level of complexity) are not considered further. This assumes the ORN receives stimulus molecules during every stimulus presentation, so does not currently take into account the lucky ORN hypothesis (Section 3.5.4.1) where low amounts of stimulus may result in sparse reception across neurons and over repeat trials.

A total of 15 ORNs, at stimulus loads of 10^{-7} g, 10^{-8} g, 10^{-9} g and blank, were examined using the UEA procedure. The above criteria were then applied to assess how many neuron-stimulus load combinations contained UEs that could imply temporal encoding of the stimulus. The number of neurons with UEs of interest, UEs that fail the criteria and no UEs at all are shown in Table 4.2. In both cases of bin sizes (5 ms in Table 4.2 a)) and 10 ms in Table 4.2 b)) there are more neurons showing UEs that fail the criteria than UEs of interest, and the number of neurons with UEs of interest increases with stimulus load. The larger bin size (10 ms) results in slightly more neurons with UEs of interest (21% of ORN-stimulus load combinations), than the smaller bin size (5 ms; 16% of ORN-stimulus load combinations). Three example neurons with UEs of interest using both sizes of time

bin are shown in Figure 4.8. The size of the time bin has little effect on the numbers of UEs found for each neuron.

Before drawing any conclusions from this section, it is important to test whether the main assumption of the UEA method holds, that of stationarity of spike trains. If spike trains are non-stationary within the period of time for which analysis takes place, the calculation of the predicted number of occurrences of each binary pattern, and hence which patterns are classed as UEs, is not a true representation of the actual data.

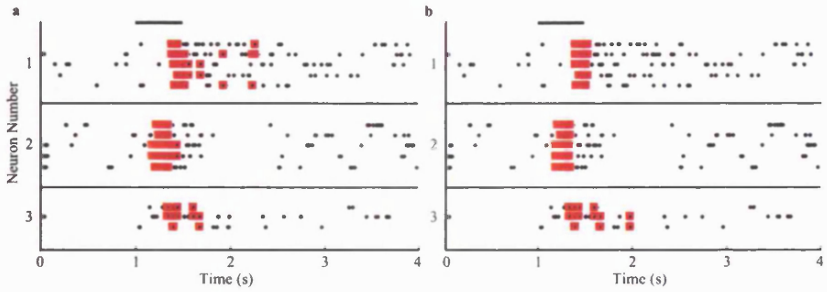

Figure 4.8: Unitary Events in ORN Recordings. Horizontal bar shows the stimulus duration, dots show time bins containing at least one spike, and red squares show spikes within UEs. a) Unitary Events with Bin Size of 5 ms. Three example ORNs show UEs in response to a stimulus load of 10^{-7} g. Each neuron shows several UEs, with multiple repeats of some patterns. All UEs occur after the stimulus onset, and comply with the criteria set out in Section 4.4.3. b) Unitary Events with Bin Size of 10 ms. Using the same three neurons (and stimulus load) as in a), UEs are again found after the stimulus onset. Neuron 3 displays an extra repeat of a UE due to the larger bin size, while neurons 1 and 2 have fewer UEs (or repeats) than in a).

4.4.4 Non-Stationarity of Recordings Breaks Fundamental Assumption of Unitary Events Analysis

The stimulus load with the highest number of neurons exhibiting UEs of interest is also the detection threshold (10^{-7} g, Table 4.2) as found in Chapter 3. The detection threshold was determined assuming a rate encoding scheme, suggesting that the firing rate of the neurons increases with the presentation of the stimulus. This implies that the spike trains are non-stationary around the stimulus onset time, which breaks one of the fundamental assumptions of the UEA method (stationarity is assumed when calculating the predicted

a) Bin size 5 ms	Stimulus loads					
	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
No UEs	6	6	5	1	18	31%
No Stim UEs	8	3	3	4	18	31%
Single UEs	1	4	3	2	10	17%
$< N/2$	0	0	1	2	3	5%
UEs of interest	0	1	2	6	9	16%
Total	15	14	14	15	58	
%	0%	7%	14%	40%	16%	

b) Bin size 10 ms	Stimulus loads					
	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
No UEs	7	5	4	1	17	29%
No Stim UEs	3	3	4	3	13	22%
Single UEs	3	4	2	2	11	19%
$< N/2$	2	2	1	0	5	9%
UEs of interest	0	0	3	9	12	21%
Total	15	14	14	15	58	
%	0%	0%	21%	60%	21%	

Table 4.2: Neurons Containing Unitary Events. UEs occurring in individual ORNs and stimulus loads were separated according to the criteria described in Section 4.4.3. Key: No UEs - no UEs at all were found; No Stim UEs - UEs were found, but these do not appear to be due to the stimulus presentation (criteria 1 and 2); Single UEs - UEs possibly relating to the stimulus were found, but only a single instance of a single UE occurred (criteria 3); $< N/2$ - several UEs (and/or instances of UEs) possibly relating to the stimulus occurred, but fewer than half of the recordings were involved (criteria 4); UEs of interest - UEs were found to comply with all criteria; % (row) indicates the percentage of neurons exhibiting UEs of interest at that stimulus load; % (column) indicates the percentage of neurons exhibiting different types of UEs across all stimulus loads. a) Breakdown of UEs Found per Neuron per Stimulus Load using a Binsize of 5 ms. As the stimulus load increases, the numbers of neurons containing no UEs, or UEs that break any of the criteria decreases. However, even at the highest stimulus load tested here, 10^{-7} g, only 40% of the ORNs exhibit UEs of interest. Over the total number of ORNs and including all stimulus loads, only 16% of possible ORN-stimulus load combinations result in UEs of interest. b) Breakdown of UEs Found per Neuron per Stimulus Load using a Binsize of 10 ms. The same general trend is observed in that as the stimulus load increases, the number of discarded UEs decreases, with the highest stimulus load again producing more UEs of interest (60%). A slightly higher proportion of ORN-stimulus load combinations result in UEs of interest (21%) for the larger bin size, but this value is low if the UEs are to represent the method by which stimulus presence is conveyed to the AL.

number of occurrences of each UE). The stationarity of each spike train was assessed using the method described in Section 4.3.6 and the results are displayed in Table 4.3. The numbers of stationary spike trains were found not to differ between the two different sizes

of time bins (5 ms and 10 ms) used for the UEA in Section 4.4.3. In general, the numbers of stationary spike trains decreased as the stimulus load increased, with the blank presentation having the largest percentage of stationary spike trains (59% for 5 ms bins and 57% for 10 ms bins). Over all the spike trains for each neuron presented with each stimulus load, only 39-40% of spike trains were stationary, suggesting that many UEs found using the UEA method (Section 4.3.4) should not be further considered. It is possible to perform the UEA method on non-stationary data by finding UEs in shorter time windows instead of the whole recording (with the firing rate throughout each window being stationary, but allowing each window to have a different rate). This is, however, not appropriate for the ORN data in its current form since the firing rate in each window is too low to allow for the statistical comparison of expected and observed coincidences (Grün et al., 2001b; Roy et al., 2000).

4.4.5 Investigation of Synchronous Spikes over Repeated Trials across Multiple Neurons

A possible solution to the non-stationarity issue discovered in Section 4.4.4 is to use UEMWA (described in Section 4.3.7). However, for each individual recording from a neuron used in the UEA method (Section 4.4.3) the UEMWA method requires further repeats from the same neuron, which are not available from the current data set. Instead, the UEMWA method can be used to search for temporal patterns of spikes occurring synchronously across the population of ORNs, for which multiple repeats are available. In total, five recordings from each ORN, at each stimulus load, with at least one spike occurring in the time interval from 1 second before stimulus onset to 3 seconds after stimulus onset were used. This resulted in different numbers of ORNs being analysed at each stimulus load (10^{-7} g used 11 ORNs; 10^{-8} g used 8 ORNs; 10^{-9} g used 10 ORNs; blank presentation used 9 ORNs). The analysis was performed using a p-value of 0.05 throughout, with both 5 ms and 10 ms time bins being tested but with little difference in outcome (representation using the 10 ms time bins are shown in Figure 4.9). The UEMWA method uses a sliding window

a) Bin size 5 ms	Stimulus loads					
	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	41	26	16	23	106	39%
Non-stationary	29	39	48	48	164	61%
Total	70	65	64	71	270	
%Stationary	59%	40%	25%	32%	39%	

b) Bin size 10 ms	Stimulus loads					
	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	40	27	18	22	107	40%
Non-stationary	30	38	46	49	163	60%
Total	70	65	64	71	27	
%Stationary	57%	42%	28%	31%	40%	

Table 4.3: Evaluation of Stationarity of Spike Trains. The stationarity of spike trains was assessed using the method described in Section 4.3.6. These tables use a filter size equivalent to 0.5 seconds (100 5 ms time bins and 50 10 ms time bins). Filter sizes equivalent to 1 second and 0.25 seconds were also tested but with little effect on the numbers of stationary spike trains (see Appendix C.2). a) Stationary of Spike Trains in 5 ms time bins. Breakdown of the stationarity of ORN spike trains according to stimulus load with the general trend of increasing stimulus load results in decreasing numbers of stationary spike trains. Over all combinations of ORNs and stimulus loads (with a maximum number of five spike trains per combination), only 39% of spike trains were deemed stationary. b) Stationarity of Spike Trains in 10 ms time bins. The effect of the bin size of the stationarity produces little difference in the number of stationary spike trains. As the stimulus load increases, the number of stationary spike trains decreases, and overall only 40% of spikes trains were stationary.

upon which UEA is performed. Initially, a sliding window of duration 50 ms was used with the aim of creating stationary recording sections, but allowing for differing rates in each window. To better assess where the UEs found occurred, a moving average was used to measure the rate of UEs along the duration of the recordings. If UEs were found relating to the stimulus presentation, the UE rate would be expected to exhibit a peak at these times. This clustering of UEs is observed for stimulus loads of 10^{-7} g and 10^{-8} g, but not for 10^{-9} g and blank presentations (Figure 4.9). Since the firing rate of some of the ORNs in response to stimulus loads of 10^{-7} g and 10^{-8} g increases, this clustering of UEs could be UEs of interest, or they could be due to windows in this time period no longer containing stationary spike rates.

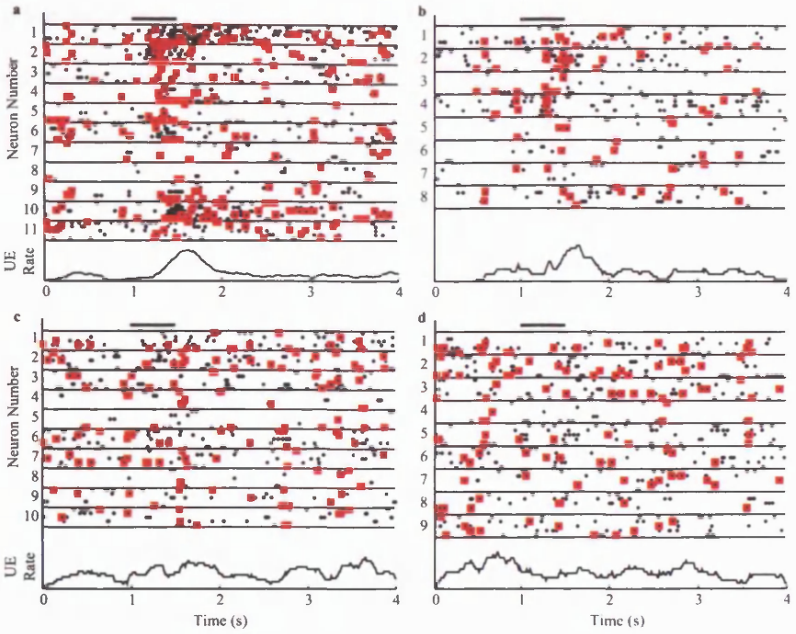

Figure 4.9: Unitary Events by Moving Window Analysis of Multiple Neurons. Short horizontal bar represents the stimulus duration, black dots indicate the spike times, and red squares indicate those spikes participating in unitary events. Five recordings from each neuron were used, with each neuron separated by a horizontal black line. Lowest panel in each subplot shows the rate at which UEs occurred smoothed using a moving average, and scaled to show the detail (scaling is different in each subplot). Spike trains were first split into 10 ms time bins and analysed using a sliding window of length 50 ms, and a p-value of 0.05. a) Stimulus Load of 10^{-7} g. 11 ORNs satisfied the criteria set out in Section 4.4.5 and UEMWA was performed. There are a large number of UEs occurring both before the stimulus onset, and after the expected response time (Section 4.4.3). There is an increase in the firing rate of the ORNs during the stimulus presentation causing a clustering of UEs (indicated by the raised UE rate). b) Stimulus Load of 10^{-8} g. 8 ORNs satisfied the relevant criteria at this stimulus load. Again, there are several UEs occurring outside the time period of interest, with a slight increase in the ORN firing rates during the stimulus presentation resulting in a clustering of UEs. c) Stimulus Load of 10^{-9} g. UEMWA was performed on 10 ORNs at this stimulus load. The stimulus presentation produces little change in the ORN firing rate, and the rate at which UEs were found remains fairly constant throughout the duration of the recordings. d) Blank Stimulus Presentation. 9 ORNs were included in the UEMWA for the blank stimulus presentation. Again, the stimulus presentation produces little change in the ORN firing rate and the UE rate remains fairly constant throughout.

Further analysis of the size of the sliding window reveals a trade-off: longer sliding windows become non-stationary, but shorter windows contain too few spikes for statistical analysis of the numbers of predicted and observed spike coincidences. This problem is examined by assessing the numbers of spikes in each set of concatenated windows (or sections) for

different window lengths and stimulus loads (Table 4.4). Up to five repeats at each stimulus load are used for each neuron, giving section durations of 250 ms, 500 ms, 2500 ms and 5000 ms for window durations of 50 ms, 100 ms, 500 ms and 1000 ms respectively. Mean numbers of spikes per section are low, especially for the shorter window lengths and lower stimulus loads (less than 3 spikes per section). According to Roy et al. (2000), the UEA method should only be used with firing rates of 10 spikes s^{-1} or higher, which translates to 2.5, 5, 25 and 50 spikes per section for window durations of 50 ms, 100 ms, 500 ms and 1000 ms. For the ORNs used in this analysis, the number of spikes per section is much lower than this (about half). The problem with low firing rates is that the frequency of occurrence of UEs becomes a random variable, making it difficult to determine whether the UEs represent a change in actual neural synchrony or are due to random variations (Roy et al., 2000). Extending the duration of the sliding window further is unlikely to remove this problem, and longer sliding windows result in sections that are no longer stationary (for a sliding window of 50 ms, 77% of sections are stationary compared to 36% of sections being stationary for a sliding window of 1000 ms).

	Stimulus loads			
Window Size	0	10^{-9} g	10^{-8} g	10^{-7} g
50 ms	0.83 (7)	0.80 (9)	0.77 (13)	1.26 (22)
100 ms	1.66 (10)	1.59 (13)	1.54 (19)	2.53 (42)
500 ms	8.38 (25)	7.93 (25)	8.00 (40)	13.01 (98)
1000 ms	16.76 (33)	15.87 (35)	16.79 (47)	27.73 (118)

Table 4.4: Spike Numbers within Concatenated Windows. Mean numbers of spikes are shown, along with the maximum number in brackets. For shorter window durations (50 ms and 100 ms), the mean average numbers of spikes in the concatenated windows (or sections, see Section 4.3.7 for description) are very low, only rising above 2 spikes per section with the window length of 100 ms and stimulus load of 10^{-7} g. Longer window lengths (500 ms and 1000 ms) have higher mean average spikes per section.

4.5 Discussion

4.5.1 Investigation into Encoding Schemes at the Periphery is Inconclusive

Detection limits at the AL cannot be accounted for by averaging the firing rates of converging ORNs alone (Chapter 3). It is possible that precise spike times in ORNs, that are not visible as a change in firing rate, could be used to pass odour information on to the AL. If this were the case, it can be assumed that for the coding scheme to be robust these spikes would occur at the same time (in relation to the stimulus) over repeated recordings from the same neuron. An initial investigation using stimulus correlation techniques suggested that although the peripheral encoding scheme is likely to be rate based, the possibility of a temporally based scheme could not be ruled out. Further investigation of the encoding scheme within individual ORNs showed little evidence of a stable, repeatable temporal pattern of spikes. The main assumption of UEA, that of spike train stationarity, was consistently violated by the ORN spike train data, resulting in unreliable occurrences of UEs.

Investigation of the encoding scheme across the population of ORNs also showed little evidence for a reproducible temporal structure, with the stationarity assumption again violated, even when analysis windows were reduced in length. This method of population stimulus encoding requires a mechanism for synchronisation across multiple ORNs. There is no evidence that ORNs across distinct sensilla are able to coordinate their activity to achieve the necessary synchronisation, so ORNs would need to be synchronised with the stimulus onset. The present study uses recordings from ORNs recorded from at different times, in different animals, resulting in possible differences in time between stimulus release from the dispenser and stimulus reception at the sensilla. More conclusive results in terms of whether a rate or temporal encoding scheme is employed could arise from the investigation of a number of simultaneously recorded ORNs.

4.5.2 Assessment of Assumptions

Each of the methods used in this chapter have some fundamental assumptions associated with them. This section examines and discusses the effects of these assumptions in the context of the results.

1. The preliminary analysis performed on the ORNs assumes that encoding and integration time windows can be defined.

   The criteria for determining whether an encoding scheme is rate based or temporally based stem from stimulus reconstruction analysis techniques and other analytical approaches (Theunissen and Miller, 1995). These methods require that an encoding time window (the duration of a spike train assumed to correspond to a single symbol in the code) and an integration time window (the duration of a stimulus that affects the response of the neuron) can be identified. The lack of temporal dynamics in the stimulus used throughout this thesis (a single 0.5 second pulse of a pheromone component) means that the encoding and temporal time windows are not possible to determine experimentally, but must be coarsely estimated. Since the results from this investigation are used only to suggest the possibility of the encoding scheme being temporal, and not to specifically state that this is truly the case, the problems associated with defining these parameters is reduced.

2. Unitary events indicative of temporal encoding occur within a limited time period after the stimulus onset.

   Spiking events prior to the stimulus onset can be classified as spontaneous or background firing. The earliest time at which a neuron can convey the presence of a stimulus is the stimulus onset. The time at which a neuron stops responding to the stimulus is variable and depends on the neuron in question (some respond for the

duration of the stimulus while others show responses long outlasting the stimulus presentation, Anderson et al., 1998; Carlsson and Hansson, 2002). Neurons in the AL, the next stage of olfactory processing, have been shown to respond within 300 ms of a stimulus (Christensen and Hildebrand, 1987; Hartlieb et al., 1997; Sadek et al., 2002) suggesting that at least some of the ORNs convey the presence of the stimulus within this time period. For the analysis in this chapter, the latest time considered for UE relating to the stimulus has been set at 1 second after the stimulus onset. UEs later than this were initially examined, but consisted of low complexity spike patterns and were unlikely to be related to the stimulus presentation.

3. ORN recordings are assumed to be stationary for the unitary events analysis.

   The statistical test in UEA compares the predicted number of occurrences of a binary pattern with the empirical number of occurrences. In order to calculate the predicted number of occurrences, the firing rate of the neuron is used to first find the probability of observing a spike in each time bin for each spike in a binary pattern. These are multiplied together to give the probability of this binary pattern occurring. The predicted number of occurrences is the probability multiplied by the number of time bins in the recording. The assumption of stationarity of spike trains is essential to accurately calculate these predicted occurrences of binary patterns, and any violation may result in an invalid statistical test producing unreliable UEs.

These were the main assumptions made whilst assessing the possibility of temporal encoding at the periphery. The first assumption, that an encoding and an integration time window can be found, is only of high importance if the analysis were used to show that a temporal encoding scheme is in use. Here, however, this method is used only to imply the possibility of a temporal encoding scheme, with UEA then used to further examine the encoding scheme. The second assumption, regarding the timing of UEs with respect to the

stimulus presentation, was used to reduce the number of spurious UEs, some of which were the result of the statistical test level of $\alpha = 0.05$, and as such is a very loose constraint. By far the most important assumption is that of stationarity. Violation of this assumption renders any UEs found to be unreliable and inconclusive.

4.5.3 Further Investigations

4.5.3.1 Further Analysis of Current ORN Data Set

If the initial investigation into the encoding scheme is assumed to be valid, i.e. there is a possibility that the ORNs employ a temporal encoding scheme; there are a variety of other techniques that could be used to further assess these spiking patterns including the following:

1. Joint Peri-Stimulus Time Histograms. The jPSTH (Aertsen et al., 1989) is a method designed to investigate the dynamics of correlation between neurons. The jPSTH matrix is built up by taking each stimulus trial and plotting the spikes of each neuron, one horizontally across the columns and one vertically up the rows. Matrix elements are incremented if both neurons have a spike at that time. This procedure is repeated for each stimulus presentation and the final matrix can be plotted in two dimensions with a colour bar representing the numbers of coincident spikes. The jPSTH allows the user to infer dependence of spike trains, gives an indication of the nature of the relationship, and displays the stimulus-related dynamics of the relationship. Statistical calculations are, however, problematic when the data is non-stationary. The jPSTH method only compares pairs of neurons, so to compare N neurons in pairs, $\frac{N!}{(N-2)!\times 2!}$ different jPSTHs would be required. The investigation of relationships between more than two neurons is currently outside the scope of this method.

2. Sliding Window Cross Correlograms. A cross-correlogram can be used to determine whether a spike in neuron A at time t affects the probability of neuron B spiking at

time $t+k$ (where k can be positive, negative or zero). The relationship is averaged over the duration of the spike trains, and no information is available to determine whether k is related to a temporal aspect of the recordings (e.g. a stimulus onset). A method of performing pair-wise cross-correlation on a sliding window has been described (Laurent and Davidowitz, 1994; Laurent et al., 1996). The cross-correlation that is calculated in each window can be represented in a matrix with the columns representing specific time lags of the cross-correlation, and the rows representing successive sliding windows. Correlations can then be interpreted according to external influences, but again there is currently no scope for the comparison of more than two neurons.

3. Cost-based Spike Metrics. Victor and Purpura (1996; 1997) propose an approach for the analysis of temporal coding using cost-based metrics. A spike train metric is basically a rule that assigns a non-negative number to pairs of spike trains expressing how dissimilar they are. A simple example would be to set a cost for inserting or removing a spike, and a cost to move a spike by time t (i.e. a cost is set for each of a number of types of transformations). The cost of the total number of transformations required to convert one spike train into another is a measure of how similar the two spike trains are.

4. Detecting Stable Spike Patterns. Fellous et al. (2004) found that the same repeated stimulus can produce more than one temporal pattern, and the pattern of spikes produced may depend on the pre-stimulus history of the neuron. These patterns are found by clustering trials according to a measure of trial similarity and reliability. The algorithm presented was developed specifically for the purpose of searching for patterns in repeated trials from the same neuron, so could be used to find a temporal encoding scheme in individual ORNs, and also across the population. A preliminary investigation here (see Appendix D) suggests that more repeats from ORNS, especially at the lower stimulus loads, would be required for a full analysis.

4.5.3.2 Further Analysis Requiring Additional Data Sets

The current data set was recorded sequentially from individual ORNs, with a different animal used for each different neuron, with a single 0.5 s odour pulse, resulting in several problems. Firstly, when the recordings are lined up with the stimulus onset, they are not necessarily lined up with when the stimulus odour molecules actually reach the antenna (due to air disturbances, for example), possibly resulting in precisely timed spikes being missed with the UEA method. Secondly, when investigating temporal spike patterns over a population of ORNs, the use of multiple animals makes success unlikely. Thirdly, the use of a single, 0.5 s odour pulse for the stimulus presentation results in a stimulus duration longer than that expected in a natural odour plume (see Chapter 6), giving rise to spike patterns perhaps not used by the moth when navigating such odour plumes. The current stimulus also renders it impossible to determine suitable encoding and integration windows for the stimulus correlation method (Lemon and Getz, 2000). To overcome these problems, further data sets would be required:

1. Simultaneous Recordings From Multiple Neurons. Using multiple impalement rigs, it is possible to record from up to four ORNs on an antenna simultaneously. Using recordings taken with this method would greatly reduce problems associated with alignment of recordings allowing for investigation of spiking patterns across several ORNs together. Both UEA and the method to find stable spike patterns (Fellous et al., 2004) are suited to this type of data set.

2. Responses to Variable Stimulus Presentations. Since the male moth usually searches for a female by navigating an odour plume (comprising of pockets of stimulus and clean air), a 0.5 s stimulus presentation is not realistic. To mimic these pockets of pheromone in an odour plume, a variable stimulus can be created. A digital stimulus controller (valve is either open or closed to allow air flow through the stimulus pipette) can be

used to create a stimulus with random pocket and inter-pocket lengths. With a more advanced analogue stimulus controller (valve can also be partially opened allowing differing amounts of air to pass, effectively altering the perceived odour concentration at the antenna) a more realistic representation of an odour plume can be created. Both of these stimulus presentations would produce temporal dynamics allowing the calculation of encoding and integration time windows used in the preliminary investigation, and would allow for the investigation of temporal patterns of ORN spikes time-locked with the stimulus.

4.6 Conclusions

The preliminary investigation of the ORN encoding scheme, based on information theoretic measures, could not rule out completely the possibility of a temporal encoding scheme. Investigations of temporal structures as reliable spike patterns across repeated recordings from individual neurons (using UEA) were impeded by the violation of the assumption of stationarity of spike trains. Using UEMWA to investigate spike patterns across a population of neurons revealed a trade off. Shorter window lengths were required to maintain the stationarity of individual moving windows, but this resulted in too few spikes per moving window upon which to perform statistical analyses. Increasing the duration of moving windows to encompass sufficient spikes violated the stationarity assumption. The present study into the encoding scheme at the periphery is inconclusive and although the biological structure of the antennae suggests that the encoding scheme is likely to be rate-based, further analysis of moth olfactory data may prove otherwise. This hypothesis, involving the encoding scheme at the periphery (Section 3.5.4.2) is not considered further here. The final hypothesis from Chapter 3 (Section 3.5.4.4), an investigation into mechanisms by which the discriminability of the system may be boosted in the AL, is reported on in the next chapter, and incorporates an investigation into the detection of synchronous spikes (Section 3.5.4.3).

Chapter 5

Modelling of Glomerular Mechanisms

5.1 Chapter Overview

The main aim of this chapter is to develop a computational model of the site of convergence of many ORNs, as seen in the olfactory system of the moth. In Chapter 3, a simple theoretical model was not able to explain how even this vast convergence of receptor inputs could achieve the boost in sensitivity observed through assessment of detection thresholds. If the convergence site were assumed to be active, rather than passive as in the theoretical model, could this explain the boost?

The initial theoretical model described in Section 3.4.6 is first converted to a computational model in which the ORN input spike trains are inhomogeneous Poisson point processes, and the output PN simply sums these input spikes. To assess the effect of more biologically realistic input spike trains and output neuron dynamics, the computational model is expanded to include firing statistics from biological ORNs, and a leaky integrate-and-fire output PN.

These modifications to the model were still unable to account for the boost in sensitivity observed in the olfactory pathway. A more complex model, involving nonlinear thresholding subunits at the site of convergence is developed, with the aim of specifically investigating the "lucky ORN" and "synchronous spikes" hypotheses from Chapter 3 (Sections 3.5.4.1

and 3.5.4.3 respectively). This nonlinear model outperformed its linear counterpart in all tests performed.

Assumptions made throughout the design process of the nonlinear subunit model are discussed, with the main outcome being that although most of the individual elements are biologically plausible, there is as yet no proof that the ideas presented here are actually used in the biological olfactory system.

5.2 Introduction

Generation of a reliable chemosensory signal at the antennal lobe requires at least a single detection event at the periphery. Yet at extremely low concentrations, such as those detected by the moth, these are likely to become relatively rare, discrete events, due to the quantum nature of molecular stimuli. Combined radiometric and electrophysiological studies in the moth *Bombyx mori* have demonstrated that a single pheromone molecule is sufficient to elicit single action potentials or bursts in pheromone sensitive ORNs (Kaissling and Priesner, 1970; Minor and Kaissling, 2003). In this case, the task of the nervous system is to identify these events in a background of spontaneous firing activity. These pheromone-elicited action potentials occur in a background of spontaneous activity from which subsequent neural processing must differentiate a valid spiking response from background firing. This is clearly limited by the degree of spontaneous activity in individual ORNs. Results from Chapter 3 support the notion that population responses of ORNs would be required to reliably discriminate discrete stimulus events at very low concentrations.

Detecting one or a few additional firing events from a single ORN at very low concentrations challenges subsequent neural processing when background activity is relatively high. The ORNs recorded for this thesis showed a wide range of spontaneous activity with relatively high mean in response to blank stimuli (1.01 - 13.31 spikes s^{-1} with a mean of 4.28 spikes s^{-1}, $n = 28$ cells). This suggests that discrete molecular detections cannot be reliably reported by single ORNs in terms of individual pheromone elicited spikes.

The problem of efficient transmission of quantal stimuli in olfaction is shared with early retinal processing in the visual system (Field et al., 2005). On a moonless night approximately one rod in 10,000 in the retina receives a single photon during its integration time (Walraven et al., 1990). Yet subsequent processing of rod responses by retinal circuitry and interneurons has been shown through behavioural and electrophysiological experiments to reliably transmit information on only a handful of photons (Barlow et al., 1971; Mastronarde, 1983). An important step in this processing is the convergence of 20-100 rod inputs onto a single bipolar cell, reminiscent of the convergence seen during early olfactory processing. This subsequent processing has been shown to reliably recover information about discrete quanta, thereby boosting signal integrity in the face of dark noise generated by rods (Baylor et al., 1984; van Rossum and Smith, 1998). Field and Rieke (2002) found that a nonlinearity in the signal transfer from rods to bipolar cells boosts the signal to noise properties in the case of quantal photon single transmission by suppressing the dark noise in nonactivated rod cells. In the case of detecting valid photon elicited sensory events from rod responses in the presence of dark noise generated during the transduction cascade, a strategy of nonlinear summation was shown to outperform a linear summation of background noise during pooling of sensory input. By applying a nonlinear thresholding of rod signals prior to sensory integration across a population, it was possible to efficiently reject dark noise from the combined response to provide an optimal readout of rod signals at light levels close to visual threshold.

The olfactory pathway may adopt a similar strategy to boost weak ORN responses for the purpose of reliable signalling. In the highly convergent architecture of early olfactory processing, subsets of ORNs with above background firing activity could be selectively boosted and/or used to suppress background noise received from the remainder of the ORN population. Hence, it is possible that at very low concentrations in which stimuli take on quantal properties, ALNs may exploit highly heterogeneous firing rates across the total

ORN population innervating a single glomerulus. In this case a linear summation strategy of signal integration would be far from optimal, since spontaneous activity is then weighted equivalently to pheromone elicited ORN response. Rather, strategies that impose a non-linear threshold will be more successful in rejecting background activity by suppressing background noise through selective boosting of above threshold responses.

In this chapter, a computational model is developed to incorporate such nonlinear strategies of sensory integration, with the aim of enabling the efficient detection of subsets of ORNs with valid responses through the boosting of these signals, but not those from spontaneously firing cells.

5.3 Data Analysis Methods

5.3.1 Synthetic Spike Train Generation

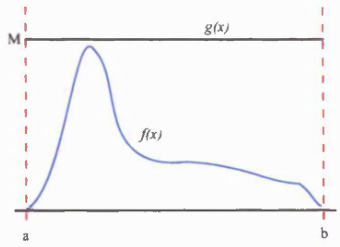

Figure 5.1: Acceptance-Rejection Method of Random Sampling. In order to sample random numbers from a distribution, $f(x)$, in the range [a, b], first find a distribution $g(x) = \mathrm{M}$ such that $\mathrm{M} < f(x)$. Generate a uniform random number, u, between [a, b], and another, v, between [0, M]. If $v < f(u)$, accept u, otherwise reject it.

To assess whether the firing rate characteristics of biological ORNs could enhance the sensitivity boost achieved by a simple model, synthetic spike trains were generated. To create these synthetic ORN spike trains, first the ISI distributions from real ORNs were created (Figure 5.6 b)). ISI's were then sampled randomly from this distribution using the acceptance-rejection method (after von Neumann, 1963; Figure 5.1):

1. Function $f(x)$ must be known over some range [a, b] (in this case, $f(x)$ is the ISI distribution), then find a second function, $g(x) = \mathrm{M}$ such that $\mathrm{M} > f(x)$.

2. Generate a random number, u, from a uniform distribution in the range [a, b].

3. Generate a random number, v, from a uniform distribution in the range [0, M].

4. If $v < f(u)$, accept random number u, otherwise, reject u and repeat steps 2–4.

Resulting synthetic spike trains have ISI characteristics of real ORNs, but with no specific spike timing information.

5.3.2 Poisson Spike Train Generation (Inhomogeneous)

For ease of control of firing rates, the model used to investigate the "lucky ORN" hypothesis (Section 5.4.3) used input ORN spike trains generated as independent inhomogeneous Poisson point processes with a different pre- and post- stimulus firing probability. In this case, the probability of observing $(N = X)$ spikes within a time period δt, is governed by the Poisson distribution:

$$P(N = X) = \frac{\lambda^X}{X!} \exp^{\lambda} \tag{5.1}$$

where λ is the mean firing rate for the neuron. Each ORN generated in this manner contained spikes with a time resolution of 1 ms, and a total duration of 1 second. Spike train durations were split exactly in half, with the first half designated as pre-stimulus, and the second half designated as post-stimulus. All ORNs had the same pre-stimulus mean firing rate, with the post-stimulus firing rate determined according to whether the ORN received the stimulus or not.

5.3.3 Poisson Spike Train Generation (Homogeneous)

Again, for ease of control of firing rates, the model used to investigate the "synchronous spikes" hypothesis (Section 5.4.4) used input ORN spike trains generated as independent homogeneous Poisson point processes. The general method is the same as that described in Section 5.3.2 and Equation 5.1, but with equal pre- and post-stimulus firing rates. Synchronous spikes were then added to those ORNs receiving the stimulus, at 500 ms $\pm j$ ms

where j is a variable amount of timing jitter. The exact time of synchronous spikes with jitter was determined by randomly sampling a value x, from a uniform distribution over the range $-j \leq x \leq j$, with the final spike time occurring at $500 + x$ ms (see Figures 5.2 and 5.3).

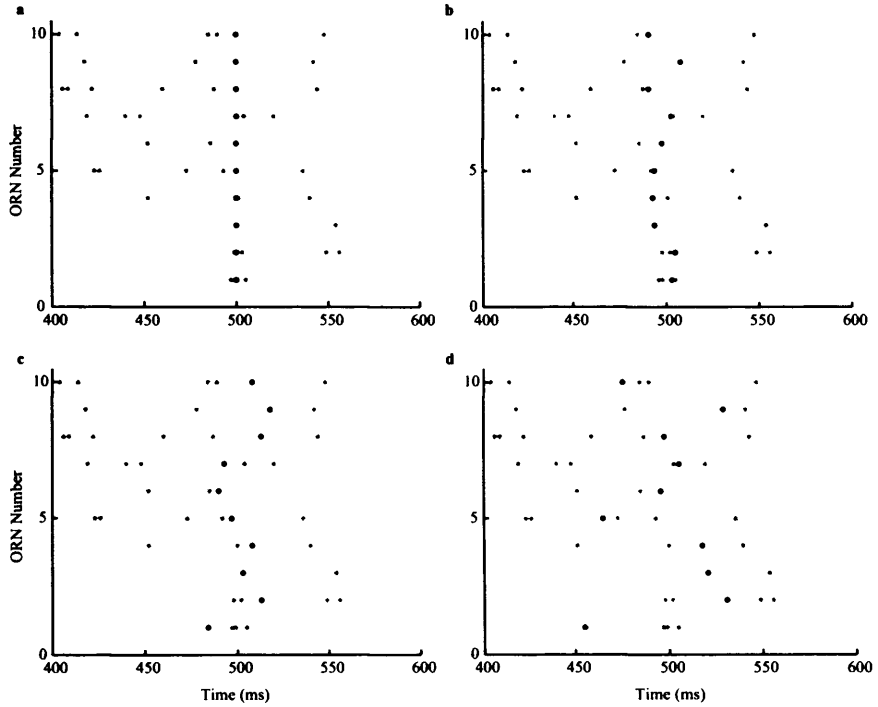

Figure 5.2: Variation of Jitter in ORN Synchronous Spike Timing. To indicate the stimulus presentation, a subset of ORNs produce a single spike with varying amounts of temporal jitter. Grey dots show the spontaneous spikes, with each ORN created as an independent homogeneous Poisson point process. Black dots show synchronous spikes representing the response to a stimulus presentation. a) Zero Jitter. All synchronous spikes occur at 500 ms. b) 10 ms Jitter. Synchronous spikes occur at 500 ms $\pm$10 ms. c) 20 ms Jitter. Synchronous spikes occur at 500 ms $\pm$20 ms. d) 50 ms Jitter. Synchronous spikes occur at 500 ms $\pm$50 ms.

5.3.4 Stimulus Presentation

When only a subset of ORNs receive the stimulus, how should these be selected? For the purpose of this chapter, two different methods are investigated. The first method, where ORNs with consecutive numbers receive the stimulus, starting from ORN_1 (Figure 5.4 a)), gets its inspiration from the natural environment of the moth. In the chemotactic search for

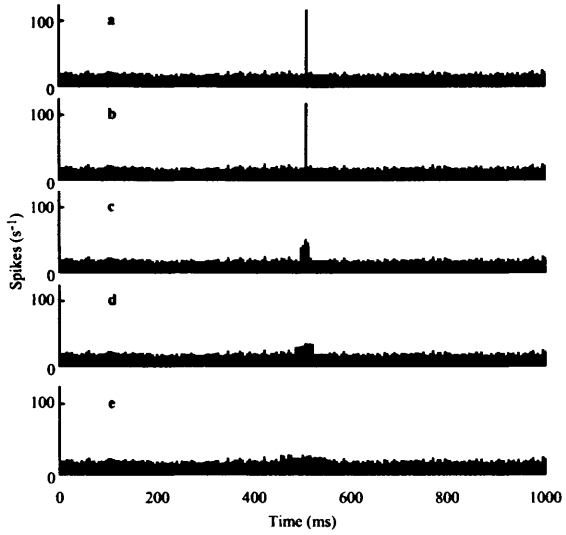

Figure 5.3: PSTHs of ORNs. Each subplot shows the spikes from 1000 ORNs split into 5 ms time bins. In each case, half of the neurons were assigned a single spike in response to a stimulus. a) Random Distribution. ORNs receiving the stimulus were assigned randomly across the population (Section 5.3.4). Stimulus spikes occurred at 500 ms, giving rise to a sharp rise in number of spikes above the background level in a single bin. b) Zero Jitter. Stimulus was assigned to ORNs 1 to 500, i.e. assigned consecutively (Section 5.3.4), with the stimulus spikes occurring at 500 ms. PSTH is the same as in a), indicating that stimulus distribution has no effect. c) 10 ms Jitter. Stimulus was assigned as in b), with the stimulus spikes occurring at 500 ms $\pm$10 ms. Rise in number of spikes is less dramatic, and is spread over 4 time bins. d) 20 ms Jitter. Stimulus was assigned as in b), with the stimulus spikes occurring at 500 ms $\pm$20 ms. The increase in number of spikes due to the stimulus is now spread over 8 time bins and is much less obvious. e) 50 ms Jitter. Stimulus was assigned as in b), with the stimulus spikes occurring at 500 ms $\pm$50 ms. Stimulus spikes are now distributed over 20 time bins and are barely discernible from the background spikes.

a female, a male moth flies upwind along a pheromone plume to locate the source (ideally a female; David et al., 1983; Vickers and Baker, 1997). The filamentous structure of the odour plume suggests that ORNs within localised areas of the antenna could encounter pheromone molecules at a particular instant in time. In *Periplaneta americana*, it has been found that some olfactory projection neurons respond only to stimuli applied to a particular part of the antenna (Hösl, 1990). A similar modular organisation has been found in the male silkmoth, *Bombyx mori* (Ai and Kanzaki, 2004). It is these ideas of structural organisation and projection patterns that gave rise to the consecutive stimulus assignment. Here, ORNs in close proximity on the antenna are represented by consecutive numbers and have the stimulus assigned to consecutively numbered ORNs.

The second method, random stimulus assignment, uses a uniform distribution to randomly select the required number of ORNs to receive the stimulus from the whole population (Figure 5.4 b)).

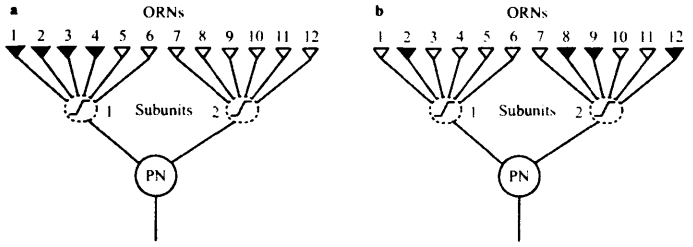

Figure 5.4: Stimulus Assignment. a) Consecutive Assignment. Four ORNs are to receive the stimulus. In this case, the stimulus is assigned to the first four consecutive ORNs. b) Random Assignment. Four ORNs are again to receive the stimulus, but in this case, the ORNs are assigned randomly from the whole population.

5.3.5 Tuning of Threshold Parameters

The model developed in this chapter uses a nonlinear threshold in each of a number of subunits, with the threshold taking the form of a sigmoid:

$$S_{out} = \frac{scale}{1 + a\exp^{-b(I_{in}-c)}} + lower \tag{5.2}$$

where *scale* and *lower* represent the maximum and minimum range of scaling factors, I_{in} and S_{out} are the input current and output scaling factor respectively, and the variables a, b, and c are used to tune the shape of the sigmoid (Figures 5.5 and 5.9 b)). The sigmoidal function acts on the value of the synaptic current at the arrival time of a spike, and provides a scaling factor by which the synaptic weight is multiplied for this spike. The values of the these parameters are described below:

1. *scale*: used to ensure a good boost to the signal. Too high a value may cause the PN to saturate. A value of $scale = 3$ is used throughout this chapter.

2. *lower*: allows a lower limit on the scaling factor. Set to $lower = 1$.

3. a: allows for fine tuning of the location of the sigmoidal function. This parameter performs a similar function to c, so is not altered in this chapter, and is set to $a = 1$.

4. b: changes the overall slope of the sigmoid. Set to $b = 20$ for this chapter, causing a sharp sigmoid (Figure 5.5).

5. c: allows for control of the location of the sigmoid, and takes a different value according to the number of ORNs in the subunit. To find the required value of c, first, a set of representative ORN spike trains were created as inhomogeneous Poisson point processes (Section 5.3.2). These were split into subunits containing the required number of ORNs, and the spike trains were summed to give a single spike train per subunit. At the time of each spike, the synaptic current was calculated using Equation 5.4 and two distributions were formed, one for the spikes that occurred pre-stimulus, and one for the spikes that occurred in response to a stimulus presentation (Figure 5.5). The value of c is chosen such that it occurs at a higher current value than most of the pre-stimulus spikes, but low enough to affect a proportion of the stimulus response spikes.

The selection of parameter values is currently experimental since there is no evidence that the olfactory system uses this type of threshold function, and so no data on which to base or test the values.

5.3.6 Receiver Operating Characteristic Curves

To assess the detection performance of the output PN from models in this chapter, ROC curves were created. The basic method was the same as that described in Section 3.3.4 and Figure 3.5, but with a slight variation in how the blank and response distributions were generated. In the case of the "lucky ORN" hypothesis, where the stimulus presentation lasted the duration of the second half of the spike train, the blank and stimulus distributions were generated by counting the number of pre- and post-stimulus spikes respectively,

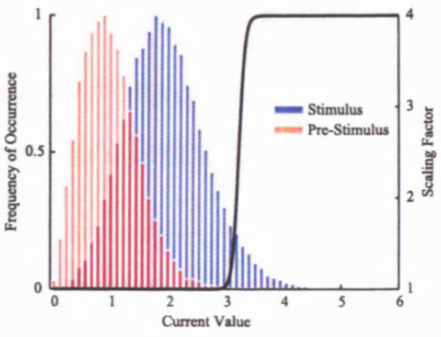

Figure 5.5: Selection of Threshold Parameters. A set of ORNs were generated as inhomogeneous Poisson point processes. These ORNs were grouped into subunits, and the synaptic current was calculated for each spike arrival time. Current values were split according to whether the spike time occurred pre-stimulus or in response to the stimulus. These distributions were plotted, and the current value at the point at which the distributions no longer overlapped was used as the parameter c. Black line indicates the sigmoidal function generated in this example, with a value of $c = 3.2$.

emitted by the PN over a number of trials. For the "synchronous spikes" hypothesis, the stimulus presentation was of a short duration, so PN spikes in response to the stimulus were counted over a period of 100 ms from the stimulus onset to generate the post-stimulus distribution. Counting the pre-stimulus PN spikes, and averaging them over a 100 ms time period generated comparable blank distributions. The control curve was generated as per Section 3.3.4, but using PN spike trains in trials where no ORNs received the stimulus.

5.4 Results

5.4.1 Convergence Models Based on Simple Summation of Inputs

5.4.1.1 Simple Model of Convergence

The simplest model to describe the convergent architecture of the early olfactory pathway is to use an output neuron that simply sums the spikes from a number of input neurons. This is based on the theoretical model developed in Section 3.4.6. A computational model, with input neurons assumed to generate spike trains following inhomogeneous Poisson point processes with a firing probability that is elevated to represent stimulus presentation, and an output neuron summing these spike trains (Figure 5.6 c)), was used to demonstrate that the discriminability (or sensitivity) enhancement of the system follows $\sqrt{n}$ (Figure 3.9 b)).

5.4.1.2 Simple Convergence Model with Synthetic Input Spike Trains

To determine whether the biological ORNs can be modelled sufficiently as Poisson processes, the coefficient of variation (CV) of the ORNs used in Chapter 3 were calculated:

$$CV = \frac{\sigma_{ISI}}{\mu_{ISI}} \tag{5.3}$$

where μ_{ISI} and σ_{ISI} are the mean and standard deviation of the interspike-intervals. CV is a measure of the within-trial variability of the neurons, and takes a value of 1 for Poisson processes regardless of the value of the mean ISI. The ORNs used in Chapter 3 had a CV lying between 0.5 and 1.5 for most values of mean ISI regardless of whether the ISIs were from pre-stimulus or during-stimulus times (blue and red dots respectively, Figure 5.6 a)). For these ORNs, the CV for the pre-stimulus ISI's lay between 0.78 and 1.69, and the CV for the during-stimulus ISI's lay between 0.61 and 1.37, which indicated that the ORNs were not Poisson. Further analysis of the ISI histograms of these ORNs showed the real neurons exhibited strong refractoriness as well as a longer tail in the ISI distribution (Figure 5.6 b)).

To incorporate these non-Poisson properties of biological ORNs into the model, the rejection method (after von Neumann, 1963, Section 5.3.1) was used to randomly sample ISIs from the real ISI distributions, with a uniform distribution as the comparison function. Resultant synthetic spike trains have ISIs, and hence firing rate characteristics of real ORNs, but don't involve any firing patterns or exact spike timing information. The output neuron again summed the input spike trains (Figure 5.6 c)). For each number of input ORNs (n; up to 1000), this convergence model was simulated 500 times to calculate the psychophysical d'. When plotted, the boost in d' again follows $\sqrt{n}$ (Figure 5.6 d)) demonstrating that non-Poisson spike train behaviour of ORNs and a simple summation representation of a readout neuron do not further enhance the discriminability of the system.

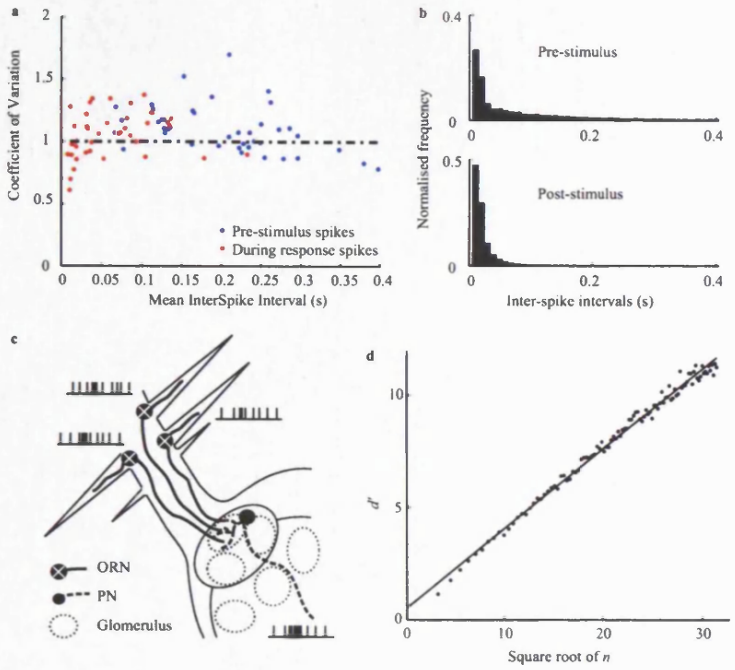

Figure 5.6: Convergence Model With Synthetic ORNs. a) Coefficient of Variation of ORNs. For a process to be Poisson, the CV should be 1 for all values of mean ISI (dashed line). For the biological neurons used in Chapter 3, ORNs display CV values ranging from 0.78 to 1.69 for pre-stimulus spikes (blue dots) and between 0.61 and 1.37 for spikes occurring during the response to the stimulus presentation (red dots), suggesting that real ORNs are not Poisson. b) InterSpike Intervals of ORNs. Pre- and post-stimulus ISI histograms show long tails in the distributions, with the pre-stimulus distribution being truncated to 0.4 s. The lack of ISIs less than 0.01 s (left most time bin) indicates the refractoriness of the neurons. c) Convergence Model. For the simpler models, ORNs are created as either independent inhomogeneous Poisson point processes, or by generating spike trains using the pre- and post-stimulus ISIs sampled from real neurons. These spike trains travel down axons to synapse onto ALN dendrites within a single glomerulus. Action potentials arriving at the synapses can be simply summed to produce the output PN spike train, or dendritic currents can be generated with exponential decay dynamics which are then integrated in the PN soma, modelled as a leaky integrate-and-fire neuron producing a spike train. d) Boost in Discriminability. A computational model using synthetic spike trains created by sampling from pre- and post-stimulus ISIs and a simple summation at the output was used to assess the boost in discriminability. For each n (up to 1000), 500 trials were performed to calculate d'. Dots show the actual values of d', with the boost in sensitivity following $\sqrt{n}$, shown by the line of best fit.

5.4.1.3 Simple Convergence Model using Integrate and Fire Output Neuron

To further discount the possibility that a simple convergence model cannot fully account for the sensitivity boost seen in the biological system, a model with more realistic signal transmission was created. Spikes from synthetic ORNs (Section 5.4.1.2) arrive at synapses,

which are modelled by generating a dendritic current with exponential decay dynamics:

$$I_{k+1} = \frac{\Delta}{\tau_e}(w\delta_{k+1,j} - I_k) + I_k \tag{5.4}$$

where Δ is the model time step, τ_e is the time constant of the exponential decay, w is the synaptic efficacy (weight), δ is the Kronecker delta and $(t_1, t_2, \ldots t_j, \ldots t_l)$ are the spike times (Destexhe et al., 1998). The dendritic current is passed into an output neuron (PN), modelled as a leaky integrate-and-fire neuron:

$$V_{k+1} = V_k\left(1 - \frac{\Delta}{\tau_m}\right) + \frac{\Delta}{C_m}I_k \tag{5.5}$$

where τ_m is the characteristic time constant for the cell, and C_m is the membrane capacitance. When the voltage reaches a fixed threshold potential, V_{th}, an output spike is emitted and the soma resets to the after-hyperpolarisation potential, V_{ahp} (Gabbiani and Koch, 1998; Mascagni and Sherman, 1998). The time constants have values in the order of tens of ms (Silver et al., 1992) with the synaptic weight w and soma capacitance C_m adjusted to ensure the output firing resembled that of a biological ALN (Figure 5.7 a)). Again, the boost follows $\sqrt{n}$ (Figure 5.7 b)) demonstrating that non-Poisson spike train behaviour of ORNs and an integrate-and-fire model of a readout neuron do not further enhance the discriminability of the system.

5.4.2 Development of Convergence Model Based on Glomerular Structure Considerations

5.4.2.1 Biological Convergence Sites

The use of more biologically realistic synapses, input firing statistics and output neuronal models has not been able to account for the sensitivity boost achieved by the male moth olfactory system when considering a simple convergence model. The site of convergence in the MGC, for a stimulus consisting of the major pheromone component, is a single glomerulus. A glomerulus itself is not strictly a single entity, but rather a densely packed area of dendrites from ORNs and ALNs containing many synapses. By staining neurons

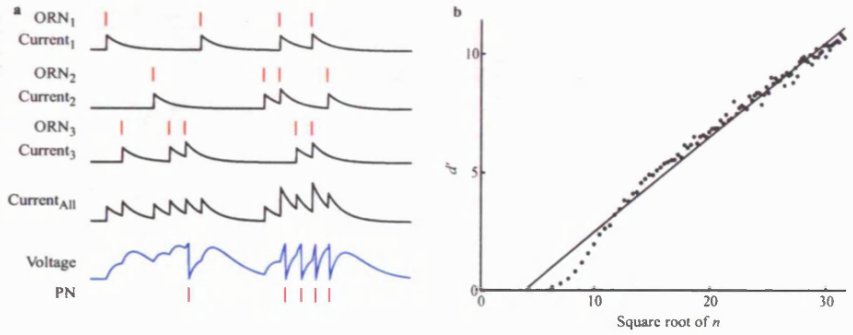

Figure 5.7: Convergence Model Incorporating an Integrate-and-Fire Output Neuron. a) Steps Within the Model. Individual spikes from input ORNs are modelled at synapses by a step increase in the synaptic current followed by exponential decay dynamics. Each neuron is modelled separately (ORN_i and $Current_i$ where $i = 1, 2, 3$) and these synaptic currents are summed at the site of convergence ($Current_{All}$). The combined current is passed into the soma of the output neuron where it is integrated giving a somatic voltage. A spike is emitted from the PN when the voltage in the soma reaches a threshold value, V_{th}. The voltage is then reset to the after-hyperpolarisation value, V_{ahp}. For numerical simplicity, the resting potential and V_{ahp} are both zero. b) Boost in Discriminability. A computational model using synthetic spike trains created by sampling from pre- and post-stimulus ISIs with exponential synaptic properties and a leaky integrate-and-fire output neuron was used to assess the boost in discriminability. For each n (up to 1000), 500 trials were performed to calculate d'. Dots show the actual values of d', with the boost in sensitivity following $\sqrt{n}$, shown by the line of best fit, indicating that further mimicking the properties of biological neurons produces no further boost in discriminability.

within antennal lobe glomeruli of the cockroach, a variety of connections between ORNs, LNs and PNs have been found (Distler and Boeckh, 1997a,b; Distler et al., 1998):

1. Monosynaptic connections from ORNs to PNs and LNs (Figure 5.8 a))
2. Polysynaptic (serial) connections from ORNs to PNs via LNs (Figure 5.8 b))
3. LN synapses onto ORNs and PNs (Figure 5.8 c))
4. Dyadic LN output synapses onto a PN and an LN (Figure 5.8 d))
5. Dyadic PN output synapses onto two LNs (Figure 5.8 e))
6. Reciprocal connections between LNs and PNs, and between two LNs (Figure 5.8 f))

This multitude of serial and reciprocal interconnections between neurons innervating the same glomerulus is likely to form the structural basis for feed-forward and feedback excitation and inhibition of neurons, and for modulation of the output activity of the glomerulus.

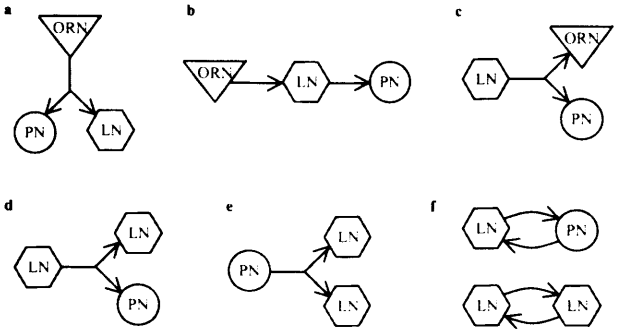

Figure 5.8: Synaptic Connections Within Glomeruli. a) ORN Synaptic Connections. ORN axons form monosynaptic contacts in a dyadic fashion onto both PN and LN dendritic processes. b) Polysynaptic (Serial) Connections. ORN axons form polysynaptic (i.e. several synapses in series) connections with the dendrites of PNs via LNs. c) LN Output Synapses I. LNs form dyadic output synapses onto ORN axons and PN dendrites. d) LN Output Synapses II. LNs form dyadic output synapses onto a PN and an LN. e) PN Output Synapses. PNs form dyadic output synapses onto two LNs. f) Reciprocal Connections. LNs form reciprocal connections to both PNs and other LNs.

Although these particular interconnections were found in the cockroach AL, there is evidence to suggest that similar synaptic connections are also likely to exist in the male moth AL (Christensen et al., 1993; Hildebrand, 1996; Homberg et al., 1989; Sun et al., 1997). Within this structure of neuropil, it is possible for far more complex mechanisms of signal transfer to occur than a simple summation of inputs.

5.4.2.2 Concept of Nonlinear Current-Based Subunits

The classical model of a neuron, where a weighted sum of synaptic inputs to the neuron is passed into a single spike generating mechanism (Section 5.4.1.3; McCullough and Pitts, 1943; Rosenblatt, 1962; Rumelhart et al., 1986) is now thought to be a poor representation of the synaptic integration in neurons with large, many-branched dendritic trees. Profusely branched dendritic trees may be capable of multiple independent nonlinear operations (Koch et al., 1982; Llinás and Nicholson, 1971; Mel, 1992a,b, 1993; Rall and Segev, 1987; Shepherd and Brayton, 1987), with anatomical, physiological and modelling studies (Häusser et al., 2000; Segev and London, 2000; Stuart et al., 1999) suggesting that some neurons may be better modelled as two-layer models (Archie and Mel, 2000; Mel et al., 1998; Poirazi et al.,

2003a,b). The first layer is made up of specific zones in the dendritic tree, which act as individual functional "subunits" performing some form of internal computation (Schiller et al., 2000; Wei et al., 2001). The second layer then operates in the manner of the traditional neuron where the subunit outputs are summed and passed through a threshold mechanism to produce the overall output of the neuron.

In the olfactory glomeruli, ORN axons synapse onto PNs, which have been shown to have large many branched dendritic arbours (Anton and Hansson, 1995; Anton and Homberg, 1999) suggesting that a better representation of the PN could be a two-layer neuronal model. In this case, the first layer consists of the ORNs forming synapses onto dendrites at the periphery of the PN dendritic tree, with individual dendritic sections acting as thresholding subunits. Outputs from these subunits are passed to the second layer, modelled as the soma of the PN, represented by a leaky integrate-and-fire-neuron (Figure 5.9 a)). The ideal action of this nonlinear signal transfer would be to reject background noise whilst boosting the signal in order to boost the sensitivity of the system. The threshold function of the subunits is represented by the sigmoid in Equation 5.2, Section 5.3.5. The sigmoidal function acts on the value of the synaptic current at the arrival time of a spike, and provides a scaling factor by which the synaptic weight is multiplied for this spike. The effect of this scaling factor on individual example ORNs is to result in an elevated synaptic current when the interval between the arrival of consecutive spikes is sufficiently small, while longer intervals result in no difference compared with the synaptic current with no scaling factor applied (Figure 5.9 c)). Synaptic currents from the individual ORNs are combined to provide the input to the second layer, the PN soma. The PN somatic voltage reaches the firing threshold more often when the input is the synaptic current with scaling factor applied to ORN spikes than without (Figure 5.9 d)).

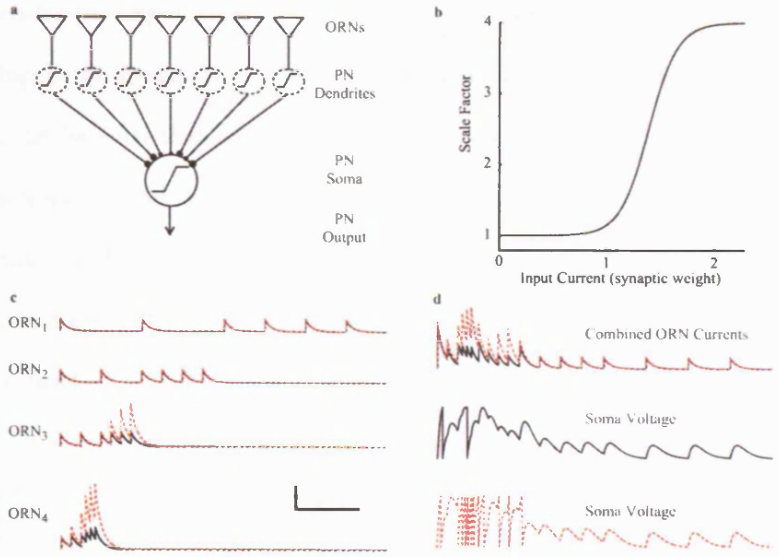

Figure 5.9: Example of Two-Layer Model with Sigmoidal Subunits. a) Schematic Representation of a Two-Layer Model. Spike trains from individual ORNs form the input to the model. These ORNs synapse onto peripheral dendritic branches of the PN, which act as individual subunits. Each of these subunits performs a sigmoidal threshold function on the synaptic weight of the input ORN spikes. Outputs from these subunits are passed onto the second layer, the PN soma, represented as an integrate-and-fire neuron. b) Sigmoidal Threshold Function. Example of a sigmoidal function (Equation 5.2) used to provide a scaling factor for the synaptic weight given an input current. c) Influence of Scaling Factor on Synaptic Currents. Horizontal scale bar: 100 ms; vertical scale bar: 2*Synaptic Weight. Black lines show the synaptic current with no modifications. Red dashed lines show the synaptic current when a scaling factor, dependent on the current and found using a sigmoidal function similar to that shown in b), is applied to the synaptic weight as each spike arrives. The spikes from four individual ORNs, represented in the dendrites of the post-synaptic neuron as exponential decays, are used to demonstrate how the effect of the scaling factor changes with interspike interval. ORN_1 and ORN_2 have interspike intervals of 125 and 62 ms, and 62 and 31 ms respectively. The synaptic current dies away within these intervals to the extent that the scaling factor on the synaptic weight induces little difference to the post-synaptic current. ORN_3 and ORN_4 have shorter interspike intervals resulting in scaling factors that produce visible increases in the post-synaptic current. d) Influence of Scaling Factor on Output Neuron. Using scale bar in c), horizontal scale bar: 100 ms; vertical scale bar: 2*Synaptic Weight for currents and 0.5*Somatic Threshold Voltage. Black lines show the synaptic current and somatic voltage with no modifications. Red dashed lines show the synaptic current and somatic voltage after the application of the scaling factor on the synaptic weight. Synaptic currents from the individual ORNs are combined on entry to the soma of the output neuron. The current produced without scaling results in 2 output spikes from the somatic voltage reaching the firing threshold. The current produced using a scaling factor on the synaptic weight results in 12 output spikes.

5.4.2.3 Introduction of Local Interneurons to Increase Output Variability

To gain most effect from sigmoidal subunits, spikes must arrive at these synapses in a timely manner. Since the duration of current input to the system lasts ~45 ms per spike (the length

of time taken for the exponential decay to reduce to less than 1% of the synaptic weight), for the scaling factor applied in the nonlinear subunits to have any effect, the post-stimulus firing rate must be such that multiple spikes are likely to arrive within this time interval. This is equivalent to spikes arriving at a synapse regularly at a rate of $\sim$22 spikes s^{-1} suggesting that the firing rates of biological ORNs are suitable for this model (mean of 5 spikes s^{-1} pre-stimulus and up to 88 spikes s^{-1} during response to stimulus depending on the stimulus load).

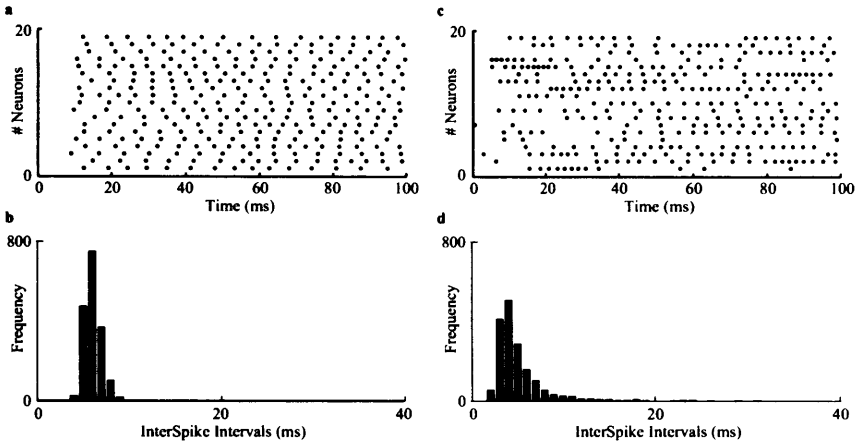

Figure 5.10: Variability of Modelled PNs. a) Modelled PN with Excitatory Inputs. When the input to the PN (modelled by an integrate-and-fire neuron), are only from excitatory ORNs (modelled as inhomogeneous Poisson point processes with a firing probability equivalent to 2.8 spikes s^{-1}), the irregularity of the input is not preserved in the output spikes, resulting in a mean CV of 0.17 (2d.p.). b) InterSpike Intervals of PN with Excitatory Inputs. The ISI histogram has a narrow range of intervals between 3 ms and 11 ms. c) Modelled PN with both Excitatory and Inhibitory Inputs. Same model as a) but with the addition of inhibitory inputs from LNs (modelled as for the ORNs, but with a negative synaptic weight). In this case, the irregularity of the input is now conserved in the output spikes, resulting in a mean CV of 1.11 (2d.p.). d) InterSpike Intervals of PN with Excitatory and Inhibitory Inputs. The ISI histogram is truncated at 40 ms, with a much wider range of intervals from 1 ms to 99 ms.

However, increasing the firing rate of the input ORNs causes problems with the output PN, even before the nonlinear subunits are incorporated. Previous studies have suggested that even with random pre-synaptic inputs, a simple integrate-and-fire neuron is not able to preserve this irregularity in its own spike output (Dayan and Abbott, 2001; Koch, 1999; Shadlen and Newsome, 1994; Softkey and Koch, 1993). This can be demonstrated using

the model developed here by using Poisson based ORNs with a pre-stimulus firing rate of 2.8 spikes s^{-1} and a simple integrate-and-fire output PN. The resulting output spike train appears to contain fairly regular interspike intervals (Figure 5.10 a) and b)), with mean CV = 0.17 (Equation 5.3). This is a much lower value than expected from biological ALNs (mean CV = 1.15). A proposed solution to this lack of output spike time variability is to incorporate inhibitory inputs to the system, providing a form of gain modulation, to balance the excitation (Calvin and Stevens, 1968; Gerstein and Mandelbrot, 1964; Shadlen and Newsome, 1994).

Within the AL of the olfactory system, inhibition arises from LN spikes, with ORNs forming excitatory input synapses onto the LNs, and the LNs forming inhibitory output synapses onto PNs (example connections 2 from Section 5.4.2.1, Figure 5.8 b)). If an ORN excites both an LN and a PN, and the LN inhibits the PN, this is known as afferent inhibition (Figure 5.11; Shepherd and Koch, 1998), and is the method by which inhibition is generated for the purpose of this model.

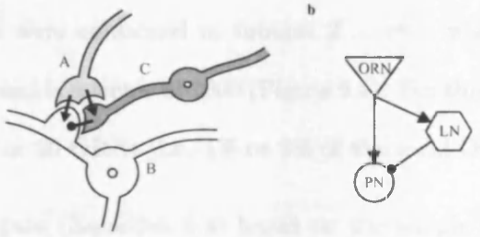

Figure 5.11: Afferent Inhibition. a) Biological Connections (adapted from Shepherd and Koch, 1998). b) Schematic Representation. ORN A excites both PN B and LN C. LN C then inhibits PN B.

Here, the synaptic currents are generated individually for the two types of neurons, with two synaptic weights, w_{ORN} taking a positive (excitatory) value for ORN spikes, and w_{LN} taking a negative (inhibitory) value for LN spikes (using Equation 5.4). These currents are passed together into the PN soma providing both excitatory and inhibitory inputs, generating a greater degree of variance in output spike times (Figure 5.10 c) and d)) with mean CV = 1.11.

5.4.3 Use of Nonlinear Current-Based Subunit Model for Lucky ORN Detection

5.4.3.1 Model Specifics

The sensitivity enhancement demonstrated in Chapter 3 gave rise to a number of possible hypotheses as to how the biological system achieves this boost. The model developed in Section 5.4.2 was used to explore one of these hypotheses, that only a "lucky" subset of ORNs from the population received stimulus molecules. In this case, the convergence model consisted of the following features:

1. One thousand input ORNs with firing rate statistics based on inhomogeneous Poisson point processes (Section 5.3.2) for flexibility.

2. One thousand inhibitory LNs with the same firing rate statistics as the ORNs.

3. Excitatory ORN synapses (Equation 5.4), grouped into subunits with a sigmoidal scaling function on the synaptic weight (Section 5.4.2.2, Equation 5.2). ORNs were grouped into subunits such that ORNs $1, 2, \ldots, n$ were contained in subunit 1, ORNs $n+1, n+2, \ldots, 2n$ were contained in subunit 2 ...etc., where n is the number of ORNs per subunit, and is a factor of 1000 (Figure 5.4). For this investigation, subunits contained either 10 or 20 ORNs (i.e. 1% or 2% of the total ORN population).

4. Inhibitory LN synapses (Equation 5.4) based on the simple summation of all input spikes.

5. Integrate-and-fire PN (Equation 5.5) with excitatory and inhibitory inputs from ORNs and LNs respectively.

For all simulations, the model was run 50 times in order to generate spike trains from 50 different PNs. To assess the performance of this model (here called the Nonlinear Subunit Model, NSM), the integrate-and-fire model, with incorporated LNs and linear synapses (i.e. using simple summation of ORN and LN inputs at the synapses, here called the Linear Synapse Model, LSM) was used as a comparison.

5.4.3.2 Detection Performance is Improved by Incorporating Nonlinear Subunits

As an initial investigation, the ORNs in the NSM were arranged into 100 subunits, containing 10 ORNs each. Input ORNs were given a pre-stimulus firing rate of 20 spikes s^{-1}, increasing to 40 spikes s^{-1} on stimulus presentation. Both these values are similar to those found in biological ORNs (see Figure 6.3 for example firing rates). The stimulus was presented to a subset of "lucky" ORNs, according to the consecutive ORNs method (Section 5.3.4), so, for example, if 10 ORNs were to receive the stimulus, these would be the ORNs numbered from 1 to 10, and would all converge onto the first subunit. Simulations were run with 0, 10, 50, 100, 200, 500 and 1000 ORNs assigned the stimulus, presented to both the LSM and NSM. The ability of the PN to detect the stimulus was assessed using ROC curves, and the corresponding AUC values (Sections 3.3.4 and 5.3.6).

For the LSM, higher numbers of ORNs receiving the stimulus allowed for reasonable discrimination from the case where no ORNs received the stimulus (Figure 5.12 a)). However, even with all ORNs receiving the stimulus, the probability of misclassification is $> 15\%$, with this accuracy decreasing as the number of stimuli ORNs decreased.

For the NSM, all numbers of ORNs receiving the stimulus resulted in higher AUC values, and so better discrimination ability, than the LSM (Figures 5.12 b) and 5.14 a)). When only 5% of the ORN population (i.e. 50 neurons) received the stimulus, the NSM shows almost perfect discrimination performance, indicated by AUC values close to 1. Comparison of the AUC values for each number of ORNs receiving the stimulus with the control curve allowed for the z-ratio to be calculated and used to identify detection thresholds (Section 3.3.4). Values of z-ratio ≥ 1.96 correspond to a significance level of $\alpha = 0.05$.

The z-ratios for the LSM suggest that when 50 ORNs or more were presented with the stimulus, the ROC curves were statistically different to the control curve (no stimulus; Table 5.1 a)). This provides an estimate of the detection threshold as 50 stimuli ORNs, so

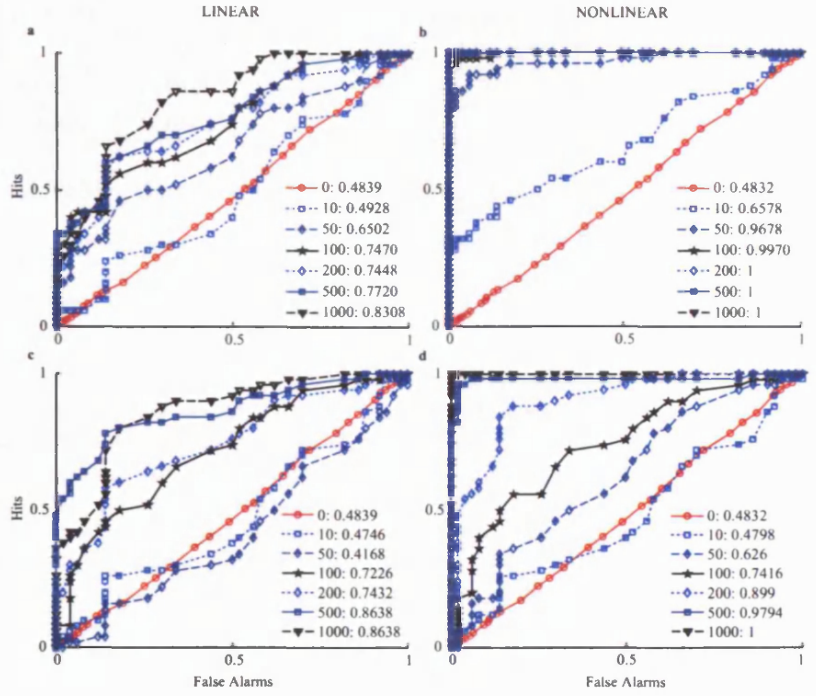

Figure 5.12: Assessment of Stimulus Assignment on Model Performances. Symbols and lines representing different numbers of ORNs that received the stimulus are preserved across the subplots. In each legend, the values are: # ORNs with stimulus : AUC value. a) and c) show ROC curves using linear synapses. In a), stimulus ORNs were assigned consecutively while in c) they were assigned randomly (see Section 5.3.4. For some numbers of stimulus ORNs, the random method gave rise to slightly higher AUC values, but for others, the consecutive method gave rise to slightly higher AUC values. On the whole, there is little difference between the two subplots suggesting that the method of stimulus assignment has little effect on the detection ability of the LSM. b) and d) show ROC curves using nonlinear subunits. In b), stimulus ORNs were assigned consecutively while in d) they were assigned randomly. b) shows better discrimination performance for all numbers of ORNs receiving the stimulus than all other subplots. d) shows similar performance to the LSM for lower stimulus numbers, but the performance is closer to that by the NSM in b) for higher stimulus numbers. This suggests that for the NSM, the method of stimulus assignment was important for lower numbers of ORNs receiving the stimulus. However, as the number of stimulus ORNs increased, and each subunit contained multiple ORNs that received the stimulus, the effect of stimulus assignment was reduced.

for the PN to be able to detect the stimulus, 5% of the ORN population must receive the stimulus. At this stimulus level, the NSM displayed almost perfect discrimination (denoted by an AUC value close to 1; Table 5.1 b)) with a probability of misclassification of $\sim 5\%$. At the lowest number of stimuli ORNs tested, the PN from the NSM was still able to detect the stimulus, suggesting a detection threshold of 1% of the ORN population (or less).

a) LSM	Control compared with number of stimulus ORNs:					
	10	40	50	100	500	1000
AUC	0.4928	0.6386	0.6502	0.7470	0.7720	0.8308
z-ratio	0.1077	1.9302	2.0837*	3.4629*	3.8572*	4.8722*

b) NSM	Control compared to number of stimulus ORNs:					
	10	40	50	100	500	1000
AUC	0.6578	0.9442	0.9678	0.9972	1.0000	1.0000
z-ratio	2.1949*	7.3450*	7.9732*	8.8210*	8.9129*	8.9129*

Table 5.1: Determination of Detection Thresholds for Linear Synapse and Nonlinear Subunit Models. Asterisks show statistically significant differences (z-ratio $\geq$ 1.96). a) LSM. A comparison of the AUC values for each number of stimulus ORNs with the AUC of the control (no stimulus) for the model with linear synapses. The z-ratio suggests that the stimulus must be presented to 50 ORNs (or more) for the ROC curve to be statistically different to that of the control curve. This gives an estimate of the detection threshold as 50 ORNs. b) NSM. Comparison of the AUC for each number of stimulus ORNs with the AUC of the control for the model with nonlinear subunits. The z-ratios suggest that all numbers of stimulus ORNs tested resulted in ROC curves statistically different from the control curve. This gives an estimate of the detection threshold as 10 ORNs.

While the detection thresholds for both the LSM and NSM were similar, the actual levels of discrimination achieved (in terms of the probability of misclassification) were far better for the NSM than the LSM.

5.4.3.3 Detection Capability is Reduced by Random Assignment of Stimulus

To investigate the effect of the arrangement of stimulus presentation, the stimulus was assigned both randomly and consecutively (Section 5.3.4) to subsets of "lucky" ORNs and presented to both the LSM and NSM. Again, subsets used were 0, 10, 50, 100, 200, 500 and 1000 ORNs, and the ability of the PN to detect the stimulus was assessed using ROC curves.

Levels of discriminability for the LSM were similar, regardless of the method of stimulus distribution (compare Figure 5.12 a) and c), consecutive and random assignment respectively). However, for the NSM, the stimulus assignment played an important role when fewer numbers of ORNs received the stimulus (Figures 5.12 b) and d) and 5.14 a)). With consecutive stimulus assignment, the PN achieved perfect discrimination when 200 or more

ORNs received the stimulus (AUC values of 1). With random stimulus assignment, 1000 stimuli ORNs still achieved this level of discrimination, but lower numbers of stimuli ORNs resulted in lower AUC values. Detection thresholds were again estimated by the calculation of z-ratios.

a) LSM	Control compared with number of stimulus ORNs:					
	10	60	70	100	500	1000
AUC	0.4746	0.6096	0.6802	0.7226	0.8638	0.8638
z-ratio	-0.1142	1.5534	2.4905*	3.0953*	5.5086*	5.5086*

b) NSM	Control compared to number of stimulus ORNs:					
	10	60	70	100	500	1000
AUC	0.4798	0.6394	0.6998	0.7416	0.9794	1.0000
z-ratio	-0.0415	1.9502	2.7751*	3.3902*	8.3002*	8.9129*

Table 5.2: Detection Ability with Random Stimulus Assignment. Asterisks show statistically significant differences (z-ratio ≥ 1.96). a) LSM. Comparison of the AUC values for each number of stimulus ORNs with the AUC of the control for the model with linear synapses. The z-ratio suggests that the stimulus must be presented to 70 ORNs (or more) for the ROC curve to be statistically different to that of the control curve. This gives an estimate of the detection threshold as 70 ORNs. b) NSM. Comparison of the AUC for each number of stimulus ORNs with the AUC of the control for the model with nonlinear subunits. The z-ratio suggests that the stimulus must be presented to 70 ORNs (or more) for the ROC curve to be statistically different to that of the control curve. This gives an estimate of the detection threshold as 70 ORNs, the same as that of the LSM.

The detection threshold for the LSM using random stimulus assignment can be estimated as 70 stimuli ORNs, indicated by the ROC curve being statistically different from the control curve (Table 5.2 a)). This suggests that the method of stimulus distribution had little or no effect on the detection ability of the PN.

The detection threshold for the NSM using random stimulus assignment was also found to be 70 stimuli ORNs (Table 5.2 b)), and so required more ORNs receiving the stimulus than with the consecutive assignment. This suggests that the detection performance of the NSM was impeded by the random stimulus assignment. For higher numbers of stimuli ORNs, where each subunit consists of multiple ORNs receiving the stimulus, the detection performance as measured by the AUC values was similar regardless of stimulus assignment (Tables 5.1 b) and 5.2 b) and Figure 5.14 a)).

5.4.3.4 Nonlinear Subunits are Better Able to Detect Small Increases in Firing Rate than Linear Synapses

Up to this point, a doubling of the pre-stimulus firing rate represented the stimulus presentations. The next sections investigate the detection performance of the LSM and NSM when the pre-stimulus rate remained the same (20 spikes s^{-1}) but with the stimulus induced firing rate set to 25 spikes s^{-1}, i.e. an increase of 25%. Simulations were run with 0, 10, 50, 100, 200, 500 and 1000 ORNs receiving the stimulus (using the consecutive assignment), with the PN detection threshold estimated using ROC curves.

With just a small increase in post-stimulus firing rate, the LSM detection performance was poor, with all numbers of stimuli ORNs resulting in AUC values close to chance levels (AUC = 0.5; Figure 5.13 a)). The NSM showed slightly better discrimination (Figure 5.14 b)), but most of the AUC values were close to chance levels, even with higher numbers of ORNs receiving the stimulus (Figure 5.13 b)).

To assess the detection thresholds, the z-ratios were again calculated from the AUC values. For the LSM, the z-ratio never reached the significance level of 1.96, so the ROC curves for different numbers of stimuli ORNs were never statistically different to that of the control curve (Table 5.3 a)). The z-ratio values for the NSM are also low, but with 1000 ORNs receiving the stimulus, the value rises above 1.96 (Table 5.3 b)), suggesting the PN was just about able to detect the increase in firing rate in response to the stimulus.

5.4.3.5 Increased Numbers of Neurons in Nonlinear Subunits Improve Stimulus Detection

The NSM investigated thus far has incorporated 100 subunits, each consisting of 10 ORNs (1% of the population). In the previous section, the NSM was only just able to detect a small increase in firing rate in response to the stimulus. This section uses a model with 20 ORNs in each of 50 subunits, with the same increase in firing rate, to assess the performance enhancement achieved. Simulations were run with 0, 10, 50, 100, 200, 500 and 1000 ORNs

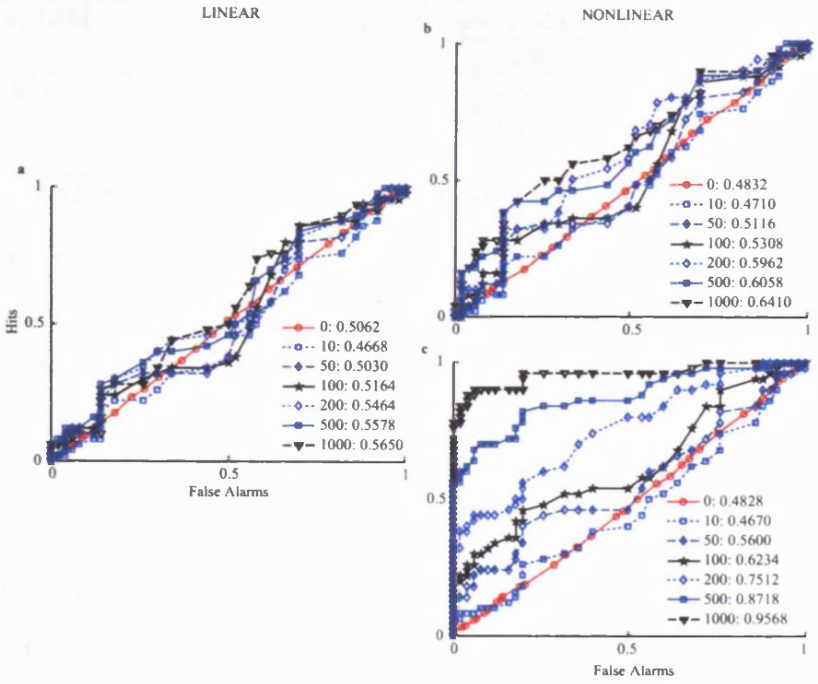

Figure 5.13: Effect of Firing Rate Increase on Model Performance. Symbols and lines representing different numbers of ORNs that received the stimulus are preserved across the subplots. In each legend, the values are: # ORNs with stimulus : AUC value. In all cases, the stimulus was assigned consecutively (Section 5.3.4). a) LSM Performance. With a small increase in firing rate in response to the stimulus, the AUC values for the LSM were close to chance levels (AUC = 0.5) for all values of stimulus ORNs. b) NSM Performance. The performance of the NSM was only slightly better than that of the LSM, indicated by increases in AUC values, particularly at higher numbers of stimulus ORNs. c) NSM with 20 ORNs per Subunit. When the NSM incorporated 20 ORNs into each of 50 subunits (instead of the original 10 ORNs into each of 100 subunits), AUC values were much higher for all numbers of stimulus ORNs.

receiving the stimulus (using the consecutive assignment), with the PN detection threshold estimated using ROC curves.

The AUC values for all numbers of ORNs receiving the stimulus were higher for the NSM with 20 ORNs per subunit (NSM20) than for the original NSM with 10 ORNs per subunit (Figure 5.13 b) and c) respectively, and Figure 5.14 b)). When 500 or more ORNs received the stimulus, the NSM20 showed a probability of misclassification of $< 15\%$, a level not achieved by the original NSM, even when all of the ORNs received the stimulus. The z-ratios calculated from the ROC curves were used to provide an estimate for the

a) LSM	Control compared with number of stimulus ORNs:					
	10	50	120	130	500	1000
AUC	0.4668	0.5030	0.5064	0.5150	0.5578	0.5650
z-ratio	-0.4815	-0.0396	0.0018	0.1066	0.6308	0.7197

b) NSM 10	Control compared to number of stimulus ORNs:					
	10	50	120	130	500	1000
AUC	0.4710	0.5116	0.5330	0.5512	0.6058	0.6410
z-ratio	0.1489	0.3462	0.6079	0.8316	1.5152	1.9713*

b) NSM 20	Control compared to number of stimulus ORNs:					
	10	50	120	130	500	1000
AUC	0.5120	0.5600	0.6608	0.6750	0.8718	0.9568
z-ratio	0.0744	0.6623	1.9491	2.1405*	5.3487*	7.3028*

Table 5.3: Determination of Detection Ability when Stimulus Presentation Produces a Small Increase in Firing Rate. Asterisks show statistically significant differences (z-ratio ≥ 1.96). a) LSM. Comparison of AUC values for different numbers of stimulus ORNs with that of the control for the linear synapse model. The z-ratio didn't reach the threshold for statistical significance for any number of stimulus ORNs suggesting that the PN was not able to detect the stimulus presentation. b) NSM with 10 ORNs per subunit. Comparison of AUC values for different numbers of stimulus ORNs with that of the control for the nonlinear subunit model with 10 ORNs per subunit. Detection performance was slightly better than that of the LSM, indicated by slightly higher AUC values. The z-ratio for 1000 stimulus ORNs suggests that the ROC curve was statistically different to that of the control curve, and hence gives an estimate of the detection threshold as 1000 ORNs. c) NSM with 20 ORNs per subunit. Comparison of AUC values for different numbers of stimulus ORNs with that of the control for the nonlinear subunit model with 20 ORNs per subunit. In this case, the AUC values were much higher than for the LSM and the original NSM. When 130 ORNs received the stimulus, the resultant ROC curve was statistically different to the control curve, providing a detection threshold estimate of 130 stimulus ORNs.

detection threshold of the NSM20 model as 130 stimuli ORNs (Table 5.3 c), z-ratio ≥ 1.96). This is much lower than that of 1000 stimuli ORNs for the original NSM suggesting that incorporating more ORNs into nonlinear subunits improved the detection performance of the PN.

5.4.4 Use of Current-Based Subunit Model for the Detection of Synchronous ORN Spikes

5.4.4.1 Model Specifics

The model developed in Section 5.4.2 was also used to explore the hypothesis that a stimulus presentation could induce a single synchronous spike in a subset of ORNs in response. In this

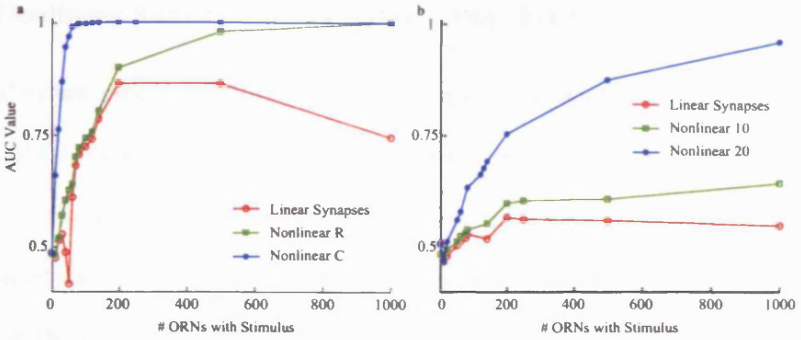

Figure 5.14: Comparison of AUC Values. Nonlinear R: NSM with random stimulus assignment; Nonlinear C: NSM with consecutive stimulus assignment; Nonlinear 10: original NSM with 10 ORNs to each of 100 subunits; Nonlinear 20: NSM with 20 ORNs to each of 50 subunits. a) AUC Values When Stimulus Response was High. AUC values are shown when different numbers of ORNs received the stimulus, when the stimulus presentation resulted in a doubling of the pre-stimulus firing rate. The NSM far outperformed the LSM when the stimulus was presented to consecutive ORNs, indicated by higher AUC values. The NSM provided near perfect discrimination (AUC = 1) when 50 or more ORNs received the stimulus. This level of discrimination was not achieved by the LSM, even when all ORNs received the stimulus. When the stimulus was assigned randomly, the performance of the NSM was degraded, particularly at lower stimulus levels where the AUC values resembled those of the LSM. b) AUC Values When Stimulus Response was Low. AUC values are shown when different numbers of ORNs received the stimulus, when the stimulus resulted in an increase of 25% of the pre-stimulus firing rate (assigned consecutively). Discrimination performance was poor for the LSM, with AUC values staying close to the chance level (AUC = 0.5), even when all ORNs received the stimulus. The original NSM (with 10 ORNs converging onto each of 100 subunits) performed slightly better, resulting in slightly higher AUC values. When the NSM was modified such that each of 50 subunits received the input spikes from 20 ORNs, the detection performance was greatly improved.

case, the convergence model consisted of the same features as in Section 5.4.3.1, but with the ORNs represented as homogeneous Poisson point processes (Section 5.3.3) allowing for the inclusion of a single response spike, with a variable amount of temporal jitter, in a subset of ORNs. While it was possible to vary the number of ORNs per nonlinear subunit, for this investigation into the detection of synchronous spikes, the number remained constant, at 10 ORNs per subunit (i.e. each subunit grouped together 1% of the input ORNs, and there were 100 subunits). The model was again run 50 times for each simulation to generate spike trains from 50 different PNs. The performance of the NSM was again assessed against the LSM.

5.4.4.2 Nonlinear Subunits Outperform Linear Synapses

As an initial investigation, ORN spike trains were generated, with the stimulus presentations to a subset of ORNs according to consecutive ORNs (see Section 5.3.4), with a timing jitter of 0 on the synchronous spikes. Simulations were run with 0, 10, 50, 100, 200, 500 and 1000 ORNs assigned the stimulus. These ORNs were presented to both the LSM and NSM with the ability of the PN to detect the stimulus assessed using ROC curves (Sections 3.3.4 and 5.3.6).

For the LSM, higher numbers of ORNs receiving the stimulus (500 and 1000) allowed for good discrimination from the case where no ORNs received the stimulus (Figure 5.15 a) and 5.16 a)). This discrimination is demonstrated by AUC values close to 1, implying probability of misclassification of $\sim$5% and $\sim$10% for 1000 and 500 ORNs respectively. Lower numbers of ORNs receiving the stimulus had much poorer detection accuracy demonstrated by ROC curves lying close to the control curve (0 stimuli ORNs), where the probability of error is close to 50%.

For the NSM, all numbers of ORNs receiving the stimulus resulted in higher AUC values, and so better discrimination ability, than the LSM (Figures 5.15 b) and 5.16 a)). When 200 ORNs received the stimulus (or more), the PN displayed perfect discrimination (AUC = 1), and when as few as 50 ORNs were assigned the stimulus, the PN still showed a probability of misclassification of less than 20%. Detection thresholds were again estimated by the calculation of z-ratios. Values of z-ratio ≥ 1.96 correspond to the significance level of $\alpha = 0.05$.

The z-ratio values for the LSM show that when 220 ORNs or more were presented with the stimulus, the ROC curves were statistically different to the control curve (no ORNs were presented with the stimulus; see Table 5.4 a)). This provides an estimate of the detection threshold as 220 stimuli ORNs i.e. for the PN to detect the stimulus, 22% of ORNs must have received the stimulus presentation. In contrast, at this stimulus level, the

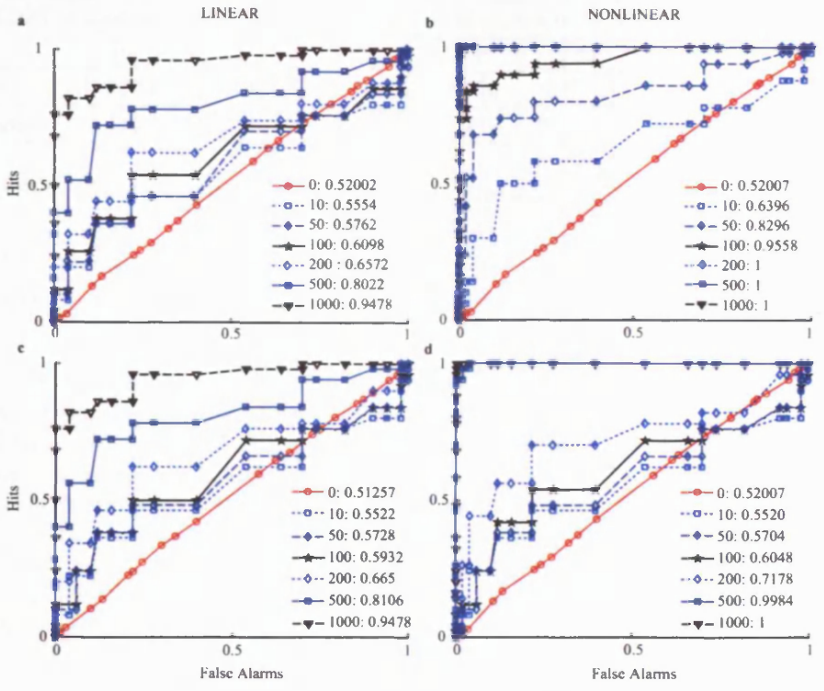

Figure 5.15: Effect of Assignment of Stimulus ORNs on PN Stimulus Discrimination. Symbols and lines representing different numbers of ORNs that received the stimulus are preserved across the subplots. In each legend, the values are: # ORNs with stimulus : AUC value. In all cases, the spikes in response to the stimulus occurred at 500 ms. a) and c) show ROC curves using linear synapses. In a), stimulus ORNs were assigned consecutively while in c) they were assigned randomly (see Section 5.3.4). There is very little difference between the two subplots suggesting that the method of stimulus assignment had little or no effect (as could be predicted from Figure 5.3 a) and b)). b) and d) show ROC curves using nonlinear subunits. In b), stimulus ORNs were assigned consecutively while in d) they were assigned randomly. b) shows better discrimination performance for all numbers of ORNs that received the stimulus compared with all other subplots. d) shows similar performance to the LSM for lower stimulus numbers, but the performance is closer to that by the NSM in b) for higher stimulus numbers. This suggests that for a model containing subunits, with lower numbers of stimulus ORNs, the method of stimulus assignment was important for the discrimination ability of the PN. For higher numbers of stimulus ORNs, the effect of stimulus assignment was much less since each subunit contained multiple ORNs that received the stimulus.

NSM showed perfect discrimination (denoted by an AUC value of 1, Table 5.4 b)). Only the lowest number of stimuli ORNs did not produce an ROC that was statistically different from the control curve. This statistical analysis gives an estimate of the detection threshold as 20 stimuli ORNs, i.e. 2% of the ORN population. Comparing the two different models, the LSM required far more ORNs to receive the stimulus to allow for detection by the PN,

a) LSM	Control compared with number of stimulus ORNs:						
	10	20	100	210	220	500	1000
AUC	0.5554	0.5550	0.6098	0.6724	0.6792	0.8022	0.9478
z-ratio	0.4329	0.4280	1.1109	1.9275	2.0196*	3.8682*	6.8500*

b) NSM	Control compared to number of stimulus ORNs:						
	10	20	100	210	220	500	1000
AUC	0.6396	0.6888	0.9558	1.0000	1.0000	1.0000	1.0000
z-ratio	1.4928	2.1504*	7.0530*	8.2795*	8.2795*	8.2795*	8.2795*

Table 5.4: Determination of Detection Ability of Linear Synapse and Nonlinear Subunit Models. Asterisks show statistically significant differences (z-ratio ≥ 1.96). a) LSM. Comparison of the AUC for each number of stimulus ORNs with the AUC of the control (no stimulus presentation) for the model with linear synapses. The z-ratio suggests that the stimulus must be presented to 220 ORNs (or more) for the ROC curve to be statistically different from the control curve, and hence for the PN to be able to detect the stimulus presentation. This gives an estimate of the detection threshold as 220 ORNs, or 22% of the population. b) NSM. Comparison of the AUC for each number of stimulus ORNs with the AUC of the control for the model with nonlinear subunits. The z-ratios suggest that all numbers of stimulus ORNs tested, apart from the lowest, resulted in ROC curves that were statistically different to the control curve. This gives an estimate of the detection threshold as 20 ORNs, or 2% of the population.

than required by the NSM.

5.4.4.3 Random Distribution of Stimulus Reduces Detection Capability

To investigate the effect of the arrangement of the stimulus presentation, the stimulus was assigned both randomly and consecutively (Section 5.3.4) to subsets of 0, 10, 50, 100, 200, 500, and 1000 ORNs. All synchronous spikes in response to the stimulus again occurred at 500 ms with no timing jitter. These ORNs were presented to both the LSM and NSM with the ability of the PN to detect the stimulus assessed using ROC curves.

For the LSM, levels of discriminability according to the number of stimuli ORNs were very similar, regardless of method of stimulus presentation (compare Figure 5.15 a) and c), and Figure 5.16 a) and b), consecutive and random assignment respectively). For the NSM, however, the method of stimulus assignment played an important role when fewer numbers of ORNs received the stimulus (compare Figure 5.15 b) and d), and Figure 5.16 a) and b) consecutive and random assignment respectively). With consecutive stimulus assignment,

the PN achieved perfect discrimination when 200, 500, and 1000 ORNs received the stimulus (AUC values of 1). With random stimulus assignment, 1000 stimuli ORNs still achieved this level of discrimination, but lower numbers of stimuli ORNs resulted in much lower AUC values, more closely resembling those achieved by the LSM (Figure 5.15 c) and d)). Detection thresholds were again determined by comparing AUC values to give z-ratios.

a) LSM	Control compared with number of stimulus ORNs:						
	20	160	170	210	220	500	1000
AUC	0.5522	0.6382	0.6596	0.6678	0.6742	0.8106	0.9478
z-ratio	0.4847	1.5678	1.8494	1.9592	2.0456*	4.1125*	6.9659*

b) NSM	Control compared to number of stimulus ORNs:						
	20	160	170	210	220	500	1000
AUC	0.5520	0.6648	0.6844	0.7294	0.7496	0.9984	1.0000
z-ratio	0.3906	1.8249	2.0900*	2.7265*	3.0274*	8.2319*	8.2795*

Table 5.5: Effect of Random Stimulus Distribution on Detection Thresholds. Asterisks show statistically significant differences (z-ratio ≥ 1.96). a) LSM. Comparison of the AUC for each number of stimulus ORNs, with random stimulus distribution, with the control AUC, for the model with linear synapses. The z-ratios and AUC values are similar to those produced with the consecutive stimulus assignment (Table 5.4 a)). For the ROC curve to be statistically different to that of the control curve, and thus giving an estimate of the detection threshold, 220 ORNs were required to receive the stimulus. b) NSM. Comparison of the AUC for each number of stimulus ORNs, using random stimulus assignment, with the control AUC of the model with nonlinear subunits. In this case, the z-ratios for the higher stimulus loads are similar to those with consecutive stimulus assignment (Table 5.4 b)). For the stimulus to be detectable by the PN, 170 ORNs must have received the stimulus presentation.

The z-ratio values for the LSM show that when 220 ORNs or more were presented with the stimulus, using random stimulus assignment, the ROC curves were statistically different to the control curve (Table 5.5 a)). This provides an estimate of the detection threshold as 220 stimuli ORNs, i.e. 22% of the population. This detection threshold is the same as that found using the consecutive stimulus assignment (Section 5.4.4.2) suggesting that for the LSM, how the stimulus was distributed across the ORN population had no effect.

For the NSM, the z-ratio values show that when 170 ORNs or more received the stimulus, with random stimulus assignment, the ROC curves were statistically different to the control

curve (Table 5.5 b)). This gives an estimate of the detection threshold as 170 stimuli ORNs i.e. 17% of the ORN population. While the NSM still required fewer ORNs to be presented with the stimulus than the LSM, the detection performance of the NSM with consecutive stimulus assignment (Section 5.4.4.2) exceeded that with the random stimulus assignment. For this NSM model arrangement, there were 100 subunits, each consisting of 10 input ORNs. For higher numbers of stimuli ORNs, when each subunit was likely to contain several stimuli ORNs, the random stimulus assignment slightly reduced the detection performance of the NSM compared to the consecutive stimulus assignment (Figure 5.16 b)). For lower numbers of stimuli ORNs, however, when each subunit contained fewer than 2 ORNs that received the stimulus (on average), the NSM performance was greatly degraded and closely resembled that of the LSM (Figure 5.16 b)).

5.4.4.4 Increase in Timing Jitter of Synchronous Spikes Degrades Detection Capability

Up to this point, all synchronous spikes produced to represent the stimulus presentation occurred at 500 ms. This section investigates the detection performance of the LSM and NSM when increasing amounts of timing jitter were added. Amounts of jitter added to the synchronous spikes started at ± 10 ms, increased to ± 20 ms, and finally to ± 50 ms (see Section 5.3.3), with the stimulus presented to consecutive ORNs. Simulations were again run with 0, 10, 50, 100, 200, 500 and 1000 stimuli ORNs, with the ability of the PN to detect the stimulus assessed using ROC curves.

For the LSM, higher numbers of ORNs that received the stimulus allowed for good discrimination of the stimulus, demonstrated by AUC values close to 1 (Figure 5.17 a), c) and e)). The probability of misclassification was $\sim$5% for 1000 stimuli ORNs, and $\sim$20% for 500 stimuli ORNs, regardless of the level of jitter introduced to the synchronous spikes (compare Figure 5.17 a), c) and e) where the temporal jitter was 0 ms, ± 10 ms and ± 20 ms respectively). Lower numbers of ORNs that received the stimulus had poor detection accuracy, and again, this was similar regardless of the level of timing jitter.

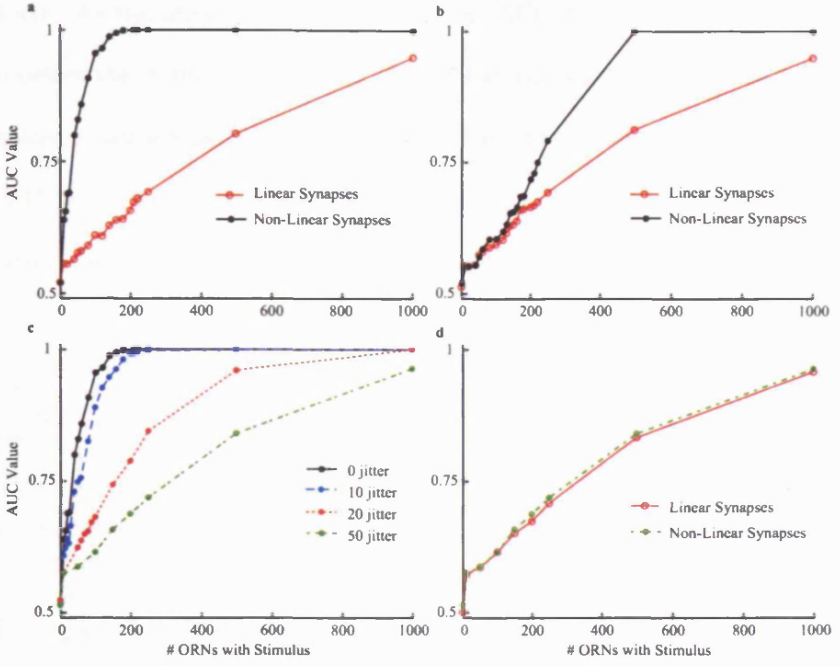

Figure 5.16: Comparison of AUC Values. a) AUC Values for Linear Synapse and Nonlinear Subunit Models. AUC values are shown, when different numbers of ORNs received the stimulus, for both the NSM and LSM. In this case, the synchronous stimulus spikes occurred at 500 ms with no temporal jitter, and the stimulus was assigned using the consecutive method (Section 5.3.4). The NSM far outperforms the LSM for all numbers of stimulus ORNs, indicated by higher AUC values. The NSM shows perfect discrimination (AUC = 1) when 200 or more ORNs received the stimulus, a level of discrimination not achieved by the LSM even when all ORNs received the stimulus. b) AUC Values with Random Stimulus Assignment. AUC values are shown, when different numbers of ORNs received the stimulus, for both the NSM and LSM. For this case, the stimulus was assigned randomly amongst the ORNs (Section 5.3.4), with no jitter on the synchronous stimulus spikes. For lower numbers of stimulus ORNs, the NSM and LSM displayed similar detection ability, but as the number of ORNs that received the stimulus rose above 200 (i.e. each subunit contained, on average, at least two stimulus ORNs), the NSM outperformed the LSM. c) Effect of Jitter on AUC Values. AUC values are shown, for different numbers of stimulus ORNs, for the NSM, with differing levels of timing jitter on the stimulus response spikes. The stimulus was assigned consecutively in all cases. As the level of timing jitter increased, the AUC values decreased, indicating a decrease in detection performance, and an increase in the probability of misclassification. d) Breakdown of Performance of Nonlinear Subunit Model. AUC values are shown, for different numbers of stimulus ORNs, with the consecutive stimulus assignment. When the level of timing jitter increased to ± 50 ms for the NSM, the detection performance was reduced to almost the level of the LSM with no jitter.

For the NSM, higher numbers of stimuli ORNs again allowed for good discrimination of the stimulus by the PN. In the case with no jitter, 200 stimuli ORNs or more resulted in perfect discrimination, indicated by AUC values of 1 (Figures 5.15 b) and 5.17 b) and

Table 5.6 a)). As the level of jitter increased, the AUC values, and hence the ability of the PN to detect the stimulus, decreased, with 500 stimuli ORNs or more, and 1000 stimuli ORNs required to achieve perfect discrimination with ±10 ms and ±20 ms jitter respectively (Figures 5.17 b), d) and f) and 5.16 c), Table 5.6 a)). With jitter levels of ±50 ms, the discrimination ability of the PN was only slightly better than the LSM with no jitter (Figure 5.16 d)).

a) AUCs	Number of stimulus ORNs:							
Jitter	10	20	30	40	90	100	160	170
0 ms	0.6396	0.6888	0.7342	0.7998	0.9390	0.9558	0.9944	0.9984
±10 ms	0.6066	0.6386	0.6628	0.7278	0.8498	0.8892	0.9622	0.9686
±20 ms	0.5756	0.5830	0.5968	0.6164	0.6704	0.6806	0.7664	0.7624
±50 ms	0.5734	0.5698	0.5752	0.5758	0.6178	0.6136	0.6550	0.6734
LSM	0.5554	0.5550	0.5646	0.5642	0.5900	0.6098	0.6388	0.6484

b) z-ratio	Control compared to number of stimulus ORNs:							
Jitter	10	20	30	40	90	100	160	170
0 ms	1.4928	2.1504*	2.7971*	3.8272*	6.6311*	7.0530*	8.1140*	8.2319*
±10 ms	1.0553	1.4651	1.7834	2.6879*	4.7001*	5.4849*	7.2012*	7.3714*
±20 ms	0.6478	0.7397	0.9122	1.1599	1.8651	2.0031*	3.2497*	3.1875*
±50 ms	0.7594	0.7148	0.7817	0.7892	1.3178	1.2642	1.8016	2.0481*
LSM	0.4329	0.4280	0.5463	0.5413	0.8617	1.1109	1.4829	1.6083

Table 5.6: Effect of Jitter on PN Detection Thresholds for the Nonlinear Subunit Model. Rows labelled LSM show the AUC and z-ratio values using linear synapses with no timing jitter. a) AUC Values. The performance of the PN in determining the presence of the stimulus increased as more ORNs received the stimulus. As the amount of jitter on the synchronous spikes increased, the performance decreased, demonstrated by lower AUC values. When the jitter reached ±50 ms, the detection performance was close to that achieved by the LSM. b) z-ratio Values. Asterisks show statistically significant differences (z-ratio ≥ 1.96). Comparisons of the AUC values for each number of stimulus ORNs with the AUC of the control curves. As suggested by the AUC values in a), the z-ratios increase as the number of ORNs that received the stimulus increased, but decreased as the level of jitter increased. Estimates of the detection thresholds at each level of jitter are: 20 ORNs for no jitter, 40 ORNs for ±10 ms jitter, 100 ORNs for ±20 ms jitter, 170 ORNs for ±50 ms jitter, and 220 ORNs for the LSM with no jitter (Table 5.4 a)).

To assess the detection thresholds of the NSM at different levels of jitter, the AUC values for different numbers of stimuli ORNs were again compared to the AUC value for

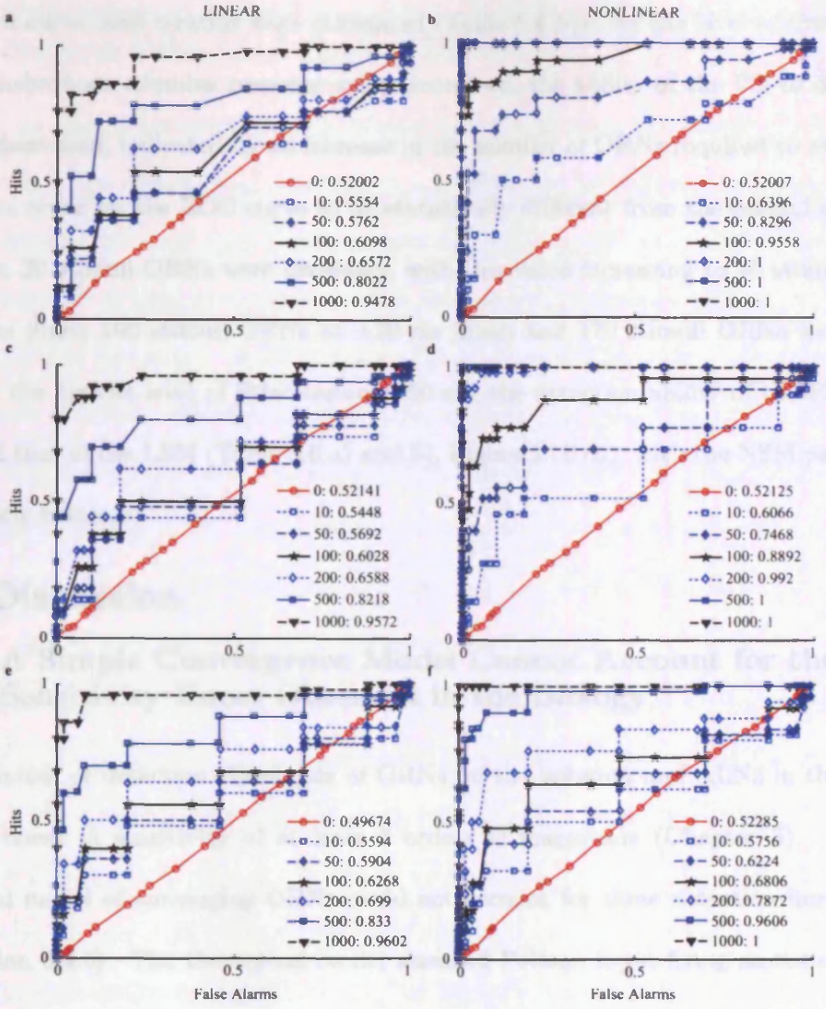

Figure 5.17: Effect of Jitter on PN Stimulus Discrimination. Symbols and lines representing different numbers of ORNs that received the stimulus are preserved across subplots. In each legend, the values are # ORNs with stimulus : AUC value. In all cases, the stimulus was assigned consecutively. Subplots on the left, a), c) and e) are the ROC curves from the linear synapse model, with no jitter, ±10 ms jitter and ±20 ms jitter respectively. There is little difference between the three subplots, suggesting that temporal jitter on the synchronous spikes had little or no effect on the ability of the PN to detect the stimulus. Subplots on the right, b), d) and f) are the ROC curves from the nonlinear subunit model, with no jitter, ±10 ms jitter and ±20 ms jitter respectively. The ability of the PN to detect the stimulus spikes, indicated by the AUC values, decreased as the amount of timing jitter on these spikes increased. This suggests that for a model incorporating nonlinear subunits, the ability to detect the stimulus relies to some extent on the precision of synchronous spikes from ORNs in response to the stimulus presentation. Even though increasing the jitter degraded the performance of the nonlinear subunit model, it still showed better discrimination for all numbers of stimulus ORNs, at all levels of temporal jitter, than the model with linear synapses.

the control curve, and z-ratios were calculated (Table 5.6 b)). As the level of timing jitter on the synchronous stimulus response spikes increased, the ability of the PN to detect the stimulus decreased, indicated by an increase in the number of ORNs required to receive the stimulus in order for the ROC curve to be statistically different from the control curve. At zero jitter, 20 stimuli ORNs were necessary, with this value increasing to 40 stimuli ORNs at ± 10 ms jitter, 100 stimuli ORNs at ± 20 ms jitter, and 170 stimuli ORNs for ± 50 ms jitter. At the highest level of jitter tested, ± 50 ms, the detection ability of the PN closely resembled that of the LSM (Table 5.6 a) and b), Figure 5.16 d)), with the NSM performing only slightly better.

5.5 Discussion

5.5.1 A Simple Convergence Model Cannot Account for the Sensitivity Boost Observed in the Biology

An assessment of detection thresholds of ORNs on the antenna and ALNs in the AL revealed a boost in sensitivity of at least 3 orders of magnitude (Chapter 3). A simple theoretical model of converging ORNs could not account for these detection limits at the AL (Section 3.4.6). The theoretical model assumed Poisson input firing statistics, with a simple summation of these spikes at a convergence site. To discount the possibility that more realistic spike trains statistics might further improve stimulus discriminability, a simple computational model was developed. Here, synthetic spike trains were generated using realistic ISI distributions taken from pheromone detecting ORNs (non-Poisson with coefficient of variation of 0.61–1.69). Unlike the Poisson spike trains considered in the theoretical model, these showed strong refractoriness as well as a longer tail in the ISI distribution. An integrate-and-fire neuron, incorporating exponential synaptic dynamics acted as the site of convergence. The boost again followed $\sqrt{n}$, demonstrating that non-Poisson spike train behaviour of ORNs and a simple model of a readout neuron do not further enhance the discriminability of the system when considering firing rates.

5.5.2 Convergence Model Incorporating Nonlinear Subunits Detects Subsets of "LuckyORNs"

Glomeruli are highly complex units innervated by several types of neurons (Distler et al., 1998). Within this structure of neuropil, it is possible for far more complex mechanisms of signal transfer to occur than a simple summation of inputs. One possible mechanism, here investigated, is that of specific zones in the dendritic arbours of PNs acting as individual functional "subunits" performing a nonlinear thresholding action on synaptic inputs from subsets of ORNs. This nonlinear subunit model (NSM) was compared with a model incorporating the same neuron characteristics as the NSM, but using only linear synapses (linear synapse model, LSM). Small subsets of ORNs were assigned a stimulus presentation to represent the "lucky ORN" hypothesis (Section 3.5.4.1), and ROC curves were used to assess the detection ability of the PNs from these two models.

When the stimulus presentation resulted in an ORN post-stimulus firing rate of double the pre-stimulus firing rate, the detection threshold of the LSM required 5% of ORNs to receive the stimulus. For the NSM, the detection threshold was 1% of ORNs. While this seems only a small improvement, ROC curves were able to show that with 5% of ORNs receiving the stimulus, the PN from the NSM displayed near perfect discrimination (with a probability of misclassification of $\sim 5\%$). This level of discrimination was not achieved by the PN from the LSM even when 100% of ORNs received the stimulus.

When the two models were given a more challenging task, an ORN firing rate increase of 25% in response to the stimulus presentation, the NSM again outperformed the LSM. For this example the PN from the LSM was not able to detect the stimulus, while the NSM detected it only with 100% of ORNs receiving the stimulus. Increasing the number of ORNs per subunit however, resulted in stimulus detection when 13% of ORNs received the stimulus.

These results suggest that if PNs were able to incorporate nonlinear subunits, subsets of "lucky ORNs" receiving the stimulus could produce a detection threshold in the AL that would be lower than that expected by measuring responses at the periphery.

5.5.3 Convergence Model Incorporating Nonlinear Subunits Detects Subsets of "Synchronous Spikes"

The same two models were used to investigate another of the possible hypotheses from Chapter 3, that of ORNs producing synchronous spikes in response to the stimulus (Section 3.5.4.3). This coding strategy, if only required to indicate the presence of a stimulus, could be achieved by each ORN transmitting a single spike, synchronised with spikes from a subset of ORNs, having little effect on the overall firing rate of the sensory neurons. Small subsets of ORNs were assigned a stimulus presentation, this time by the addition of a single spike (with varying amounts of timing jitter) to represent the "synchronous spikes" hypothesis, and ROC curves were again used to assess the detection ability of the PNs from these two models.

When the stimulus presentation resulted in a single response spike, exactly synchronised in time across a subset of ORNs, the detection threshold of the LSM PN required 22% of ORNs to receive the stimulus. The PN from the NSM required only 2% of ORNs to receive the stimulus, and showed perfect discrimination at the detection threshold of the LSM. When timing jitter was added to the synchronous spikes, as the level increased, the performance of the NSM was gradually degraded, requiring 4%, 10% and 17% of ORNs to receive the stimulus with jitter of ± 10 ms, ± 20 ms and ± 50 ms respectively.

These results suggest that a PN consisting of nonlinear dendritic subunits is easily capable of detecting a small number of coincident events at the periphery, resulting in a detection threshold much lower than expected by examining the firing rate detection threshold of the ORNs.

5.5.4 Assessment of Biological Relevance

Throughout the development of this nonlinear subunit model, inspiration came from several different areas. This section looks briefly at the different elements making up the complete model.

1. Synaptic Connections Provide a Pathway for Inhibition.

   Staining of neurons in the glomeruli of the cockroach revealed a multitude of serial and reciprocal interconnections between ORNs, LNs and PNs (Distler and Boeckh, 1997a,b; Distler et al., 1998). These connections suggest that LNs, as well as ORNs, are likely to play a role in the transmission of signals from olfactory stimuli. In the NSM, the idea of afferent inhibition (Shepherd and Koch, 1998) was used to introduce inhibitory synapses onto the PN, causing the model output firing statistics to closely resemble those of biological ALNs.

2. Dendritic Branches Can Act as Individual Functional Subunits.

   Studies of neurons with profusely branched dendritic trees suggested that some neurons might be better modelled as two-layer models (Archie and Mel, 2000; Mel et al., 1998; Poirazi et al., 2003a,b), with the first layer as specific zones in the dendritic tree acting as functional "subunits", and the second layer as a more traditional spike generating threshold mechanism. Mel (1993) and Poirazi et al. (2003a; 2003b) suggest that in the pyramidal neurons of the rat hippocampus, these individual subunits perform a nonlinear thresholding function on the synaptic inputs to dendritic branches. Since PNs often have large, many branched dendritic trees (Anton and Hansson, 1995; Anton and Homberg, 1999), it is possible that segments of dendritic branches act as thresholding subunits, and have here been modelled as such.

3. Groups of ORNs Synapse onto a Specific Area of a Glomerulus.

   In the cockroach, a number of olfactory projection neurons were found that responded only when a stimulus was applied to a specific area of the antenna (Hösl, 1990), with a similar modular organisation found in the silkmoth (Ai and Kanzaki, 2004). It follows then that individual ORNs located in close proximity on the antenna may converge onto a specific section of a glomerulus, and so preserving, in some sense, the

structural arrangement of the antenna. In the model, this was represented as ORNs with consecutive numbers converging onto the same subunit.

4. Threshold Function can be Suitably Located.

   Of the various possible parameters used to shape and position the sigmoidal threshold function in the subunits, the most important seems to be where to locate it in terms of the current coming into the threshold function (Figure 5.5). An experimental and modelling study performed at the front end of the visual system of the mouse (Field and Rieke, 2002) suggested that locating a nonlinear threshold such that all noise was rejected, more than compensated for the elimination of some stimulus responses. The threshold for the PN subunits was therefore located such that the nonlinear threshold had no boosting effect on pre-stimulus spikes (and noise from ORNs that didn't receive the stimulus). While this also meant that some stimulus responses were also unaffected, other responses were such that the signal was boosted sufficiently for detection.

5. Assignment of Stimulus.

   In view of the fact that a model of the convergence site of ORNs has been developed here, the distribution of the stimulus presentation, as well as the type and level of stimulus response of the ORNs, was a function of the model. The question of which ORNs should receive the stimulus presentation was approached in two different ways. The simplest approach was to randomly assign the stimulus to the required number of ORNs. However, this method was perhaps not well suited to the biological system. The actual stimulus presentation in terms of an odour filament in a plume, would be likely to interact with a localised area of antenna, and as such would affect a group ORNs in close proximity. The second approach incorporates this idea by assigning the stimulus to consecutively numbered ORNs in the model.

While all the features incorporated in the NSM can be found in some form in the biology, they come from diverse areas of the central nervous system, from retinal processing in the mouse, to pyramidal neurons in the rat hippocampus. There is as yet no evidence that these computational structures exist in the moth olfactory system, but it is still possible to use the model developed here to investigate possible boosting strategies, with the aim to then explore the biological system in search of the relevant mechanisms.

5.6 Conclusions

Theoretical considerations of signal convergence suggested that the observed biological boost in sensitivity (Chapter 3) is beyond that achievable by summation of the elevated firing rates over a population of receptor neurons. Since it is likely that a glomerulus is actually an active site of convergence of receptor input, an enhanced computational model has been developed, incorporating dendritic subunits with nonlinear thresholds. This nonlinear subunit model was used to investigate two of the hypotheses from Chapter 3, the "lucky ORN" hypothesis and the "synchronous spikes" hypothesis. In both cases, the discrimination performance of the nonlinear subunit model exceeded that of a comparable model, incorporating only linear synapses. This suggests that if biological glomeruli are active convergence sites, the ALNs would be able to detect lower stimulus amounts than expected by examining the firing rate detection thresholds of ORNs at the periphery.

Chapter 6

Frequency Analysis of Olfactory Receptor Neuron Responses

6.1 Chapter Overview

The main aim of this chapter is to perform an initial investigation into the temporal aspects of the ORN response to a stimulus. Chapter 3 describes the vast convergence of these neurons on arrival at the AL, presumably resulting in a high level of redundancy. Do the individual neuronal responses provide more information about the stimulus than just the chemical identity and/or stimulus load? Is it possible that these ORNs are also capable of capturing information on the frequency of interception of odour molecules?

Many repeat electrophysiological recordings from several different ORNs were taken in response to the major pheromone component in order to characterise the responses in both the time and frequency domains. Spike train timings were first extracted as detailed in Chapter 2. To perform the characterisation in the time domain, repeat recordings from individual ORNs were pooled to create a single PSTH for each ORN. The temporal responses illustrated by the PSTHs were characterised by fitting an exponential decay and calculating the decay constant for each neuron. Characterisation in the frequency domain was performed by taking fast Fourier transforms (FFT) for each PSTH to identify the frequency dynamics of each neurons response, with individual ORNs showing response characteristics in different frequency bands.

In order to compare the frequency dynamics of ORN responses with possible stimulus input frequencies, an artificial plume was generated in a wind tunnel using a wind speed of $\sim$0.56 ms^{-1} (navigable by a moth performing chemotactic search behaviour). A hotwire anemometer was used to measure local flow movements at eight different locations in the plume. The FFT was taken of the averaged anemometer signals to identify the frequencies that occur in a plume at this velocity. When compared, the frequency dynamics of ORNs coincide with those frequency components found in the wind tunnel plume.

Assumptions made throughout the investigation are discussed, with the main outcome being that the stimulus used doesn't adequately mimic the odour filaments a moth would encounter in a plume. A more realistic stimulus might provide even more interesting temporal and frequency responses in ORNs. Finally, possible biologically plausible reasons for why different ORNs appear to be able to code for different frequencies of odour filament interception are suggested and discussed.

6.2 Introduction

According to the so-called efficient coding hypothesis (Barlow, 1961, 2001), sensory systems strive to achieve an efficient representation of the multitude of different stimuli in an animal's natural environment. Two strategies are thought to be involved. Firstly, single sensory neurons should fully utilize their output capacity when encoding information. This has been demonstrated by precisely matching the tuning curve of a single neuron to the statistics of natural visual stimuli (Laughlin, 1981). Secondly, sensory neurons should provide a statistically independent response from one another. For example, recent investigations in the visual system have revealed that the responses of retinal ganglion cells to natural stimuli are relatively independent (Nirenberg et al., 2001). Both these strategies refer to the idea of natural stimuli, but how should this be defined in the case of an olfactory stimulus?

For the purpose of this chapter, the effects of natural stimuli of the olfactory system of the male moth during his search for a reproductive partner are of interest. The olfactory

environment of the moth includes odours from plants as well as odours from other animals and different species of moth. From this cacophony of odour molecules, transported in a shifting odour plume, the male moth must first identify the correct combinations of pheromone components from a calling female and then navigate the plume to find her.

Added to the problem of odour identification is the issue of plume structure. Odour molecules are not delivered to the moth antennae in a continuous stream, but in short bursts interspersed with clean air. Obstacles in the path of the plume, and the flow dynamics themselves, create turbulence, which forms this intermittent filamentous plume structure. Surrogate molecules representing the presence of odour molecules can be used to measure the structure of these turbulent plumes (using ionised air, Murlis and Jones, 1981; Murlis et al., 1992, 2000; using a tracer gas, Justus et al., 2002a, 2005). These studies have demonstrated the filamentous structure of the plume, and shown how the intermittency varies with distance from the source.

The intermittency of a plume (or the time in which the signal is or is not present) is related to the frequency at which the filaments arrive at a detector. This frequency can be measured and has been shown to vary according to the plume substance, location and velocity (e.g. an air plume of $\sim$4 m s^{-1} on a flat coastal strip contains frequencies up to $\sim$4 Hz (Murlis and Jones, 1981) whilst an air plume of $\sim$5 m s^{-1} in a wind tunnel contains frequencies up to $\sim$10 Hz (Justus et al., 2002a) and an artificial aquatic plume flowing at $\sim$0.1 m s^{-1} contains frequencies up to $\sim$2 Hz (Moore and Atema, 1991)). Given the variation in flow velocities and associated plume frequencies it is expected that different animals encounter plumes with different time and frequency characteristics in their natural environments.

In a plume, there is a range of frequencies at which odour filaments arrive at a stationary point, but is this information actually utilised by the animal? It seems that the answer to this question is dependent on species. Behavioural experiments of moths in a wind tunnel

have been used to assess whether the presence, and indeed frequency of filaments in the plume influence the upwind flight pattern. Studies have shown that some male moths are able to navigate in a homogeneous "cloud" or continuous plume of pheromone (Justus and Cardé, 2002), while this environment completely halts the upwind aspect of the flight path for other moth species (Baker et al., 1985; Justus and Cardé, 2002; Justus et al., 2002b; Mafra-Neto and Cardé, 1994, 1995a,b; Willis and Baker, 1984). These studies show that, when in a pulsed pheromone plume, the upwind flight of these moths is resumed, with higher pulse frequencies eliciting faster and straighter upwind trajectories (Justus et al., 2002b; Mafra-Neto and Cardé, 1994, 1995b). It seems that for many moth species, not only is it the filamentous structure of the plume that is important for female search strategies, but also the frequency at which filaments interact with the antennae.

For behavioural responses to differ according to the frequency at which an odour plume is pulsed (Baker et al., 1985; Justus et al., 2002b; Mafra-Neto and Cardé, 1994, 1995b), it is possible that the neurons at the front end of the olfactory system are able to resolve these filaments. Individual neurons in the AL and populations of neurons on the antennae have been investigated to determine up to what frequency pulsed stimuli can be resolved. Pheromone sensitive projection neurons have been shown to be able to resolve artificial odour pulses of up to 10 Hz (Christensen and Hildebrand, 1988, 1997). Fourier analyses of electroantennogram (EAG) recordings have shown that by measuring the whole antenna, it is possible to resolve the temporal structure of odour filaments of up to 33 Hz (Bau et al., 2002). At present, there have been very few experiments looking at individual ORNs, but recordings taken from single sensilla suggest that, depending on the receptor type/pheromone component, ORNs can follow pulsed stimuli up to 5 Hz (Kaissling, 1986a; Kodadová, 1996; Marion-Poll and Tobin, 1992; Rumbo and Kaissling, 1989).

Are individual ORNs capable of coding for different environmental factors such as plume frequency dynamics, along with the odour stimulus? By collecting a large number of electrophysiological recordings from a number of olfactory receptor neurons, it was possible to

assess the temporal structure of the neuronal response to a pheromone component. By converting these responses into the frequency domain, it has been possible to conduct an initial investigation into whether the frequency dynamics of the ORNs are compatible with those frequencies found in a natural plume. To mimic a natural plume, the frequencies found in an artificial plume generated in a wind tunnel, were compared with the ORN responses, with the finding that these frequencies overlap. When the responses of individual ORNs were examined, it was found that even with just a single pheromone pulse, different ORNs show different temporal characteristics. The response from each neuron contains dynamics at different frequency bands, suggesting the possibility that different ORNs may be able to detect different frequencies in a plume, and so providing an independent property as suggested by the efficient coding hypothesis.

6.3 Data Analysis Methods

6.3.1 Electrophysiological Recordings

Many recordings from different ORNs, found in the antenna of the male *Spodoptera littoralis*, were taken by collaborators in Sweden (see Chapter 2) in order to assess their frequency response. The stimulus used was the major component of the female sex pheromone, Z9, E11-14:OAc. Since this is vital for sexual reproduction in this species, ORNs tuned to this component are found in plentiful supply in the base of long sensilla (Ljungberg et al., 1993; Ochieng' et al., 1995). Although the detection threshold is thought to lie below 10^{-7} g per 10^{-5} l of solvent (on filter paper, delivered using a 0.4 ml s^{-1} pulse of air, Ljungberg et al., 1993; see Chapter 2), a stimulus load of 10^{-5} g, with presentation lasting 0.5 seconds was used to ensure a strong response in each recording and neuron. Between 12 and 59 recordings (with 2 seconds pre-stimulus, 0.5 seconds stimulus and ~2.5 seconds post-stimulus), were taken from 11 different ORNs responding to the presented odour. Spikes timings were extracted from each recording according to the methods described in Chapter 2.

6.3.2 Fitting of Exponential Decays

An initial characterisation of different temporal responses exhibited by ORNs was performed. To this end, exponential decays of the form $y = y_0 e^{-\lambda t}$ were fitted to the PSTHs for each neuron to describe the decay of the response (see Figure 6.3 for examples). In this case, y is the number of spikes s^{-1} recording^{-1} at time t, y_0 is the spike s^{-1} recording^{-1} at time $t = 0$, and λ is the decay constant. For this calculation, the start time of the decay was set to be the PSTH time bin containing the maximum number of spikes s^{-1} recording^{-1}, and the end time was calculated using the threshold method described in Section 3.3.2, Figure 3.1, with no time restriction on this location. The curves were fitted using MATLAB's `cftool` function in the Curve Fitting Toolbox. This uses a least squares method to minimise the summed square of residuals when obtaining the coefficient estimates (see Draper and Smith, 1998). The time taken for the response to decay to half the original spikes s^{-1} recording^{-1}, $T_{\frac{1}{2}}$, i.e. when $y = \frac{y_0}{2}$ can be calculated from the decay constant λ since rearranging the equation for the exponential decay gives $T_{\frac{1}{2}} = \frac{ln(2)}{\lambda}$.

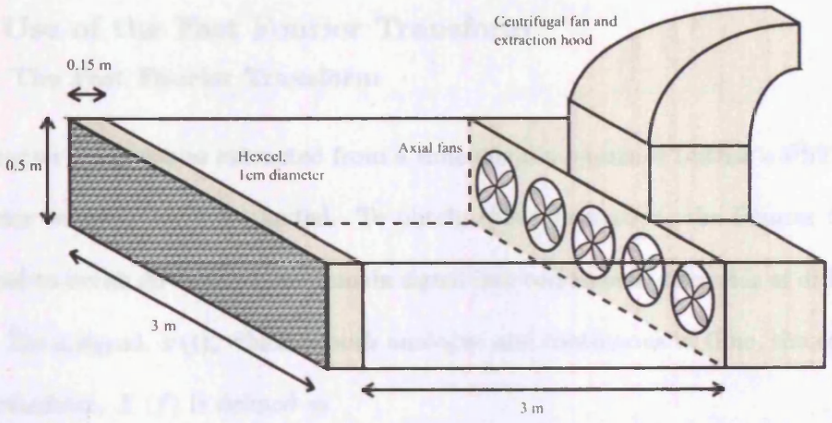

Figure 6.1: Wind Tunnel Setup. A transparent polyethylene wind tunnel, of dimensions 3 m long by 3 m wide by 0.5 m high, was used to measure flow properties of an artificial plume. The intake end was covered with 0.15 m of aluminium honeycomb (Hexcel, 1 cm diameter). Air was drawn through the wind tunnel by five axial fans located at the outflow end. Air was then expelled using a centrifugal fan and passed through an extraction hood containing charcoal filters. Adjusting the speed of the axial fans allows for different flow velocities to be created inside the wind tunnel.

6.3.3 Wind Tunnel Measurements

All wind tunnel measurements were taken by Jing Gu.

To compare the frequency dynamics of the ORNs with the frequency components contained in an odour plume, an artificial plume was generated in a wind tunnel. The wind tunnel consisted of a centrifugal fan and five axial fans used to draw air along a transparent polyethylene tunnel (3 m long by 3 m wide by 0.5 m high). An aluminium honeycomb (Hexcel, 1 cm diameter) was placed at the intake end of the wind tunnel to minimize external flow disturbances. Adjustment of individual fan speeds allowed for the control of airflow within the wind tunnel (see Figure 6.1). For this investigation, a wind speed of $\sim$0.56 ms^{-1} was used to generate a plume that would be navigable by a moth performing chemotactic search behaviour (Vickers et al., 2001; Willis and Baker, 1984). A hotwire anemometer was used to measure local flow movements at this velocity. The anemometer recordings were taken at eight different locations in the wind tunnel, and the anemometer output, measured in Volts, was sampled at 10 kHz.

6.3.4 Use of the Fast Fourier Transform

6.3.4.1 The Fast Fourier Transform

In this chapter, frequencies extracted from a time domain sequence (either a PSTH, or the anemometer output) are investigated. To obtain this information, the Fourier transform can be used to break down the time domain signal into constituent sinusoids of different frequencies. For a signal, $x(t)$, which is both analogue and continuous in time, the continuous Fourier transform, $X(f)$ is defined as

$$X(f) = \int_{-\infty}^{\infty} x(t) \exp^{(-j2\pi ft)} dt \tag{6.1}$$

where $x(t)$ and $X(f)$ are complex functions of the continuous-time variable t and the continuous-frequency variable f respectively.

For sampled data, such as that obtained throughout in this chapter, Equation 6.1 is not suitable and a discrete version must instead be used. Assuming samples are equidistant in time, the discrete Fourier transform (DFT) is given by

$$X(k) = \sum_{n=0}^{N-1} x(n) \exp^{(-j2\pi nk/N)} \tag{6.2}$$

where $k = 0, 1, 2, \ldots, N-1$. The fast Fourier transform (FFT) used in this chapter is an efficient algorithm for computing the DFT (Cooley and Tukey, 1965).

6.3.4.2 Of PSTH

To assess the frequency components of the neuronal responses, only the time period of the PSTH in which each neuron shows a response to the stimulus was used. This time period is calculated using the same threshold method as that used to find the net spikes integration time period (Section 3.3.2, Figure 3.1), but with no time restriction on the location of the end time bin. The individual PSTH segments were first multiplied by a hamming window, defined as

$$w(n) = 0.54 + 0.46 \cos\left(\frac{2\pi}{N} n\right) \tag{6.3}$$

where $n = 0, 1, 2, \ldots, N-1$ and N is the number of samples covered by the window. The purpose of the hamming window was to reduce the effects of discontinuities at the edges of the segments (Harris, 1978; Kay and Marple, 1981, Figure 6.2), and then the FFT was used to find the frequency components (Kay and Marple, 1981 and references therein).

6.3.4.3 Of Anemometer Recordings

The anemometer was placed in the wind tunnel at eight different locations. Data were sampled at each of these locations, for a time of 10 seconds, with six repeats, giving a total of 1 minute at each recording location. These signals were averaged across the locations, and the FFT taken to give the frequencies components in the plume.

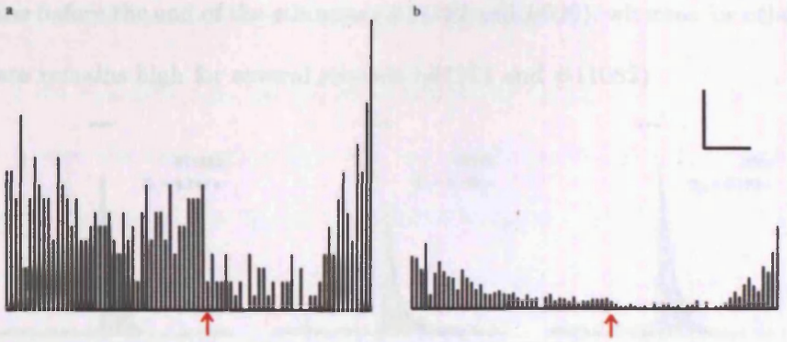

Figure 6.2: Application of Hamming Window to PSTH. Horizontal scale bar: 0.1 s, vertical scale bar: 10 spikes s^{-1} recording^{-1}. The Fourier transform assumes a periodic signal of infinite duration. For the purposes of this chapter, only segments of PSTH showing response to the stimulus are of interest, and these are neither infinite nor periodic. For this reason, the PSTH segments were multiplied by a hamming window to ensure a smooth transition between the end of a segment, and the start of the next repeat of the segment (shown by the red arrows in both subplots). a) shows the join between the end of a PSTH segment and the start of the next repeat of the segment with no pre-processing. The general trend of the data at the end of the segment is that there are more spikes s^{-1} recording^{-1} in each PSTH time bin than at the start of the segment. This results in a jump at the transition point. b) shows the same join between PSTH segments after multiplication with the hamming window. The numbers of spikes s^{-1} recording^{-1} in each bin have been suppressed close to the join, resulting in a much smoother transition from one PSTH segment to the next.

6.4 Results

6.4.1 Peri-Stimulus Time Histograms Reveal Differing Temporal Structure in Neuronal Responses

To assess the temporal structure of the neuronal responses, PSTHs were constructed. All of the recordings from each neuron were aligned with the start of the stimulus, spike timings were split into 10 ms time bins, and each time bin was corrected to display spikes per second per recording (grey bar charts, Figure 6.3). Each subplot contains the PSTH for each different neuron, with the horizontal bar representing the stimulus presentation. For each neuron, there is a short delay between the stimulus onset and a change in firing rate, caused mainly by the time taken for the odour molecules to reach the antennae. All neurons show a sharp increase in firing rate caused in response to the stimulus, although some neurons respond to a much higher degree than others (compare, for example, #11193 with #1118). In many cases, the neurons appear to act as "onset" detectors since the response

starts to wane before the end of the stimulus (#11182 and #920), whereas for other neurons, the firing rate remains high for several seconds (#1111 and #11082).

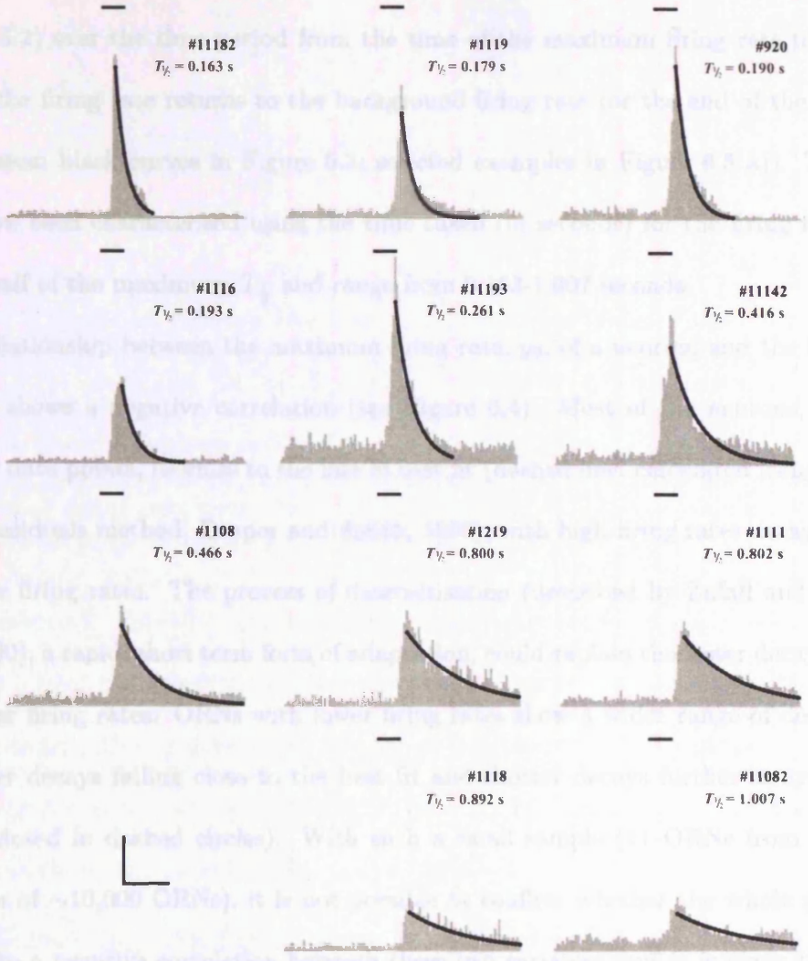

Figure 6.3: Olfactory Receptor Neurons Responses to a Pheromone Stimulus. Spikes from between 12 and 59 electrophysiological recordings for each of 11 different ORNs in response to a pheromone stimulus presentation were put into 10 ms time bins to generate the PSTHs shown in gray (horizontal black bars in each subplot show the stimulus presentation; horizontal scale bar represents 1 second; vertical scale bar represents 50 spikes per second per recording). All of the neurons show an initial sharp rise in the firing rate corresponding to the start of the stimulus presentation, but show differing characteristics as the response reduces to the background firing rate. These decay properties can be described by fitting exponential decays (see Section 6.3.2), shown as the black curves in each subplot. The time required for the response firing rate to decay to half of the maximum firing rate can be calculated from the exponential decay, and is shown as $T_{\frac{1}{2}}$ for each neuron.

6.4.2 Exponential Decays Characterise Neuronal Responses

The decay of the neuronal response can be described by fitting an exponential decay (see Section 6.3.2) over the time period from the time of the maximum firing rate to the time at which the firing rate returns to the background firing rate (or the end of the recording in some cases; black curves in Figure 6.3; selected examples in Figure 6.5 a)). The decay curves have been characterised using the time taken (in seconds) for the firing rate to die down to half of the maximum, $T_{\frac{1}{2}}$ and range from 0.163-1.007 seconds.

The relationship between the maximum firing rate, y_0, of a neuron, and the half decay time, $T_{\frac{1}{2}}$, shows a negative correlation (see Figure 6.4). Most of the neurons, shown as individual data points, lie close to the line of best fit (dashed line, calculated using the least absolute residuals method, Draper and Smith, 1998), with high firing rates decaying faster than lower firing rates. The process of desensitisation (described by Zufall and Leinders-Zufall, 2000), a rapid, short term form of adaptation, could explain the faster decay of ORNs with higher firing rates. ORNs with lower firing rates show a wider range of decay times, with longer decays falling close to the best fit and shorter decays further away (see data points enclosed in dashed circles). With such a small sample (11 ORNs from a possible population of ~10,000 ORNs), it is not possible to confirm whether the whole population conforms to a negative correlation between these two variables, but it is likely that this is not the case. As a general hypothesis, it could be expected that the line of the negative correlation forms a soft upper boundary, with high firing ORNs taking mainly shorter decay times, and lower firing ORNs taking a wide range of decay times.

6.4.3 Frequency Analysis of Neuronal Responses Suggests a Range of Frequencies of Interest

The variation in the decay rate of the different neurons suggests that the neurons themselves may be able to detect odour filaments arriving at different frequencies. To assess the frequency components, only the time period of the PSTH in which the neuron shows a

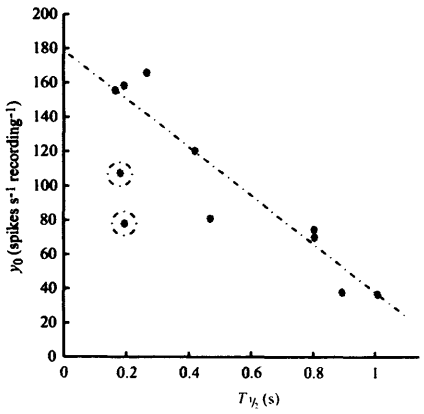

Figure 6.4: Correlation Between Neuronal Firing Rate and Decay Time. The maximum firing rate for each neuron (y_0, as calculated using the fitted exponential decays) was plotted against $T_{\frac{1}{2}}$, the time taken for the response to decay to half of the maximum firing rate. The resulting plot shows a negative correlation, with higher firing rates decaying more quickly than lower rates, demonstrated by a shorter $T_{\frac{1}{2}}$. Although many of the data points (representing individual ORNs) fall close to the line of best fit (dashed line), several points do not (two of which are surrounded by a dashed circle). These two neurons in particular have a much faster decay than other neurons with similar firing rates.

response was used. The FFT of this section was taken (see Section 6.3.4.2), and by plotting the magnitude of the transform against the frequency, it is possible to see that different neurons contain different power over a range of frequencies (selected neurons, Figure 6.5 b)). The maximum frequencies of interest range from 2-10 Hz across the different neurons (neurons not shown lay between these two extremes).

6.4.4 Frequency Analysis of an Artificial Plume Indicates Frequencies Found in Natural Plumes

Using the wind tunnel set up described in Section 6.3.3, local flow movements were measured with a hotwire anemometer. The anemometer recordings (measured in Volts), taken from different locations within the plume, illustrate changes in the filamentous structure over time and with location (see Figure 6.6 a)). Voltage levels close to zero indicate very low, or zero flow, while higher values indicate more air movement. The filamentous structure seen in Figure 6.6 a) resembles that seen in other wind tunnel experiments (Justus et al., 2002a; Murlis and Jones, 1981; Murlis et al., 2000; Vickers et al., 2001). An FFT was again used

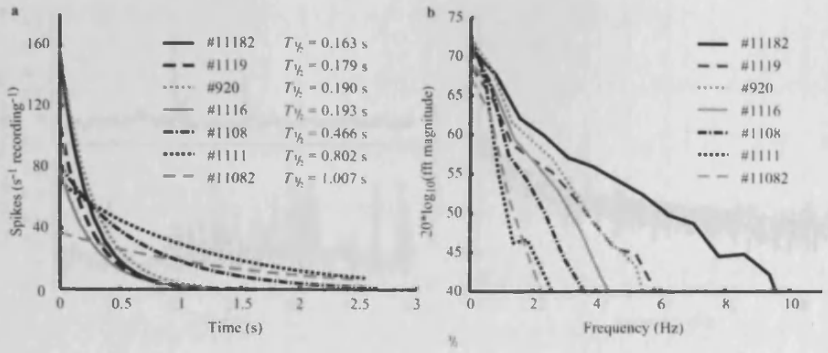

Figure 6.5: Responses of Selected Neurons. a) Exponential Decays. A selection of the fitted exponential decays demonstrates the range of times over which different neurons take for the response to decay. The half-decay-time, $T_{\frac{1}{2}}$, ranges from 0.163-1.007 seconds. b) Frequency Analysis of PSTH Shapes. FFT's were taken on the section of each PSTH with elevated firing in response to the stimulus. The magnitude of the power at each frequency up to 11 Hz is shown for each of the selected neurons shown in a). The frequencies of interest range between 2-10 Hz for the different neurons. The same line styles are used for each neuron in both subplots.

to investigate the frequency components in the plume. The averaged anemometer signal from 8 locations within the plume, each lasting for 60 seconds, was considered (see Section 6.3.4.3). In general, frequencies of interest occur before the power in the frequency spectrum flattens out, suggesting that a high proportion of the total energy in the signal from this wind tunnel generated plume is contained below 10 Hz (Figure 6.6 b))

6.5 Discussion

6.5.1 Comparison of Neuronal and Artificial Plume Frequency Spectra

Assessment of the frequency components contained within ORN responses, found by transforming the temporal structure of the ORN response to a pheromone stimulus pulse with an FFT, suggests that ORN responses contain dynamics up to 10 Hz, which may be capable of resolving corresponding pulses of the same frequency. Individual ORN responses contain very different power in their frequency spectrum over a range of frequencies, with the frequencies of interest ranging from 2-10 Hz.

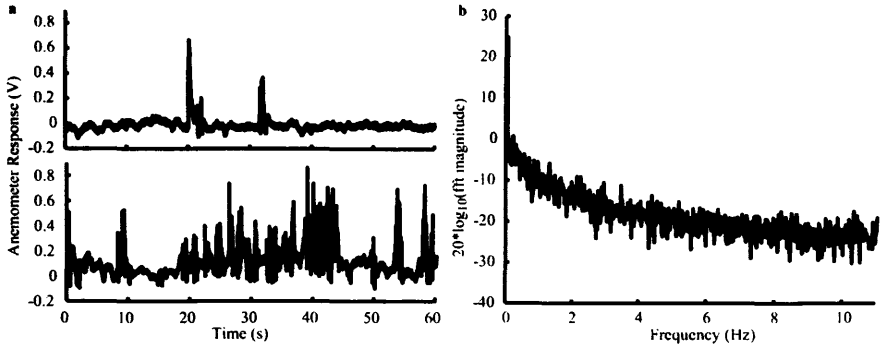

Figure 6.6: Wind Tunnel Plume Investigation. a) Anemometer Recordings Inside a Wind Tunnel. A hotwire anemometer was used to measure the localised air movements under low flow conditions (suitable for moth flight, see Section 6.3.3). The anemometer responses (measured in Volts and sampled at 10 kHz) were recorded from 8 different locations inside the wind tunnel. The responses at two different locations are shown. Levels close to zero show very little air movement whereas higher values indicate more rapid changes in local flow velocity. b) Frequency Analysis of Wind Tunnel Airflow. The recordings from each wind tunnel location were averaged. The FFT of this averaged signal was taken, and the magnitude plotted over the frequency scale used for Figure 6.5 b). The power in the spectrum decreases and flattens out after ~6-8 Hz, suggesting that these frequency components dominate the chemical signal that a moth would be likely to encounter in a pheromone plume.

Do these frequencies of interest to the moth's receptor neurons exhibit any relation to the frequencies observed in naturally occurring odour plumes? To answer this question, the frequency spectrum of an artificial plume, generated in a wind tunnel, was calculated. For this plume, the power in the spectrum flattens out at ~8 Hz, suggesting that frequencies of interest lie below this value. When compared, the frequencies found in the wind tunnel plume, and the frequencies found in the ORN temporal responses, are found to both cover a similar range of frequencies. This suggests the possibility that individual ORNs themselves may be able to capture specific bands of frequency information found within naturally occurring odour plumes.

6.5.2 Assessment of Assumptions

A number of assumptions have been made in this chapter, involving both plume and odour dynamics, and also relating to the linear analysis of ORN responses. These assumptions are stated here, with an assessment of how they may affect the overall findings.

1. An artificially created plume mimics a naturally occurring plume.

   As mentioned in Section 6.3.3, a wind speed of 0.56 ms^{-1} was used to generate a plume comparable to natural plumes navigable by a moth (Vickers et al., 2001; Willis and Baker, 1984). The use of honeycomb at the intake end, and axial fans to draw the air along the wind tunnel produce a filamentous structure with low levels of disturbance (Justus et al., 2002a,b, 2005; Mafra-Neto and Cardé, 1995a). This artificially generated plume shows structural components that are similar to those described in natural environments (compare Figure 6.6 a) with those found in Murlis, 1986; Murlis and Jones, 1981; Murlis et al., 2000).

2. An anemometer measurement of airflow represents the structure of odour filaments.

   The output voltage of the anemometer used in this chapter is sampled at 10 kHz. Since the frequencies of interest in this generated plume have been shown to lie in the range 0-10 Hz, the sampling rate used is far higher than required by the Nyquist criterion where a sampling rate of 20 Hz would be required to completely determine the plume signal by its samples (Nyquist, 1928; Shannon, 1949). The anemometer recordings themselves provide an accurate measurement of the flow of air over a hot wire. Given that the odour molecules are advected by moving air, and that the anemometer recordings show similar structure to surrogate plumes (Justus et al., 2005), a measure of the air movements is a reasonable representation of odour movements.

3. A single odour pulse represents the odour filaments found in a plume.

   Odour filaments have been measured arriving at a detector at frequencies of up to $\sim$10 Hz in a plume (Justus et al., 2002a; Vickers et al., 2001). During olfactory search, moths often surge upwind, resulting in a higher rate of interception of odour filaments (Baker and Haynes, 1989; Justus et al., 2002a). The stimulus used for this chapter was a single square pulse of pheromone, lasting for 0.5 s. This is much longer than the

estimated length of time that a moth spends in an odour filament (for example, given an interception frequency of up to 5 Hz, each filament is a maximum of 0.2 s). A single odour pulse is not a good representation of filaments in an odour plume, but ORNs still respond to the stimulus and show response dynamics up to ~10 Hz. If a more realistic stimulus were presented, ORN responses may contain dynamics at different frequencies than those shown here.

4. *Spodoptera littoralis* requires a flickering signal for olfactory search.

   Spodoptera littoralis, the Egyptian cotton leaf worm moth, found mainly in Europe and Africa, is a pest on cotton plants and a variety of vegetable crops (Tamaki and Yushima, 1974). This suggests that the search for a reproductive partner is likely to take place predominantly in fields and open spaces. Unlike the moth, *Cadra cautella*, which is a pest of stored products often in enclosed, wind-free environments, and appears to not need a flickering signal (Justus and Cardé, 2002), it is likely that *Spodoptera littoralis*, flying in open environments, requires the flickering, filamentous structure of natural plumes (Murlis et al., 2000) for olfactory search.

5. Taking the Fourier transform of a PSTH provides a good representation of the frequency dynamics of a neuron.

   The process of creating a PSTH for each neuron involves dividing the spike trains into individual discrete bins of duration Δt, and summing the numbers of spikes occurring in each bin. This procedure generates an estimate of the firing rate of the neuron, which is dependent on both the size of the time bin, and their placement. To avoid the arbitrariness of this bin placement, a window of duration Δt could be slid along the spike train with numbers of spikes within the window counted at each location. Both of these methods of estimating the firing rate perform a linear filtering of the spike trains. While this may be acceptable for the initial investigation presented here,

the underlying behaviour of neurons is not truly linear, and the result is that the FFT then shows the frequency dynamics of a linear approximation of the neural response to the stimulus.

These were the main assumptions made whilst assessing the temporal response structure of ORNs. After considering these assumptions, it is likely that the male moth requires a flickering signal, which provides the ORNs with different frequencies of stimulus arrival. The frequencies found in an artificial plume here are similar to those found in naturally occurring odour plumes, although the stimulus used here is not a good representation of these frequencies. At present, the frequency analysis is performed on a linear approximation of the neuronal response (assumption 5). If a better representation of the odour filaments in a plume was used to stimulate the moth antennae (assumption 3), or a white noise stimulus was used, a more detailed approximation of these neurons may become available, providing further insight into the frequency dynamics of these neurons. Further experimentation and a larger subset of ORNs could clarify the presence of further neurons with independent coding of frequency information.

6.5.3 Further Hypotheses

6.5.3.1 Species Determination of the Calling Female

The females of several moth species are known to rhythmically extrude their pheromone glands during pheromone release (Cardé and Roelofs, 1973; Cardé et al., 1984; Conner et al., 1980; Valeur et al., 1999). Do the females of different species release their pheromone at different pulse frequencies? There is evidence that the pulsation of the pheromone gland is required to increase the emission rate of the pheromone (Schal and Cardé, 1985), and is not distinctive for each species. Could the pulsatile release of pheromone be conserved in the odour plume and provide the male moth with additional information on the species of the calling female? This is doubtful since the structure of the pulsatile pheromone release is unlikely to be conserved in natural odour plumes due to the plume interaction with objects

in the environment creating eddies and causing turbulence. For these reasons, it seems that the purpose of different ORNs showing an interest in different frequencies in an odour plume to aid in the identification of the species of the calling female is unlikely.

6.5.3.2 Determination of Location in Odour Plume

Since ORNs show frequency dynamics of up to ~10 Hz, if they were capable of resolving odour filaments at these frequencies, could this provide information to the moth on its current position in a natural odour plume? The plume position can be thought of in two ways, firstly, the longitudinal distance from the source, and secondly, the cross-sectional position in the plume.

1. Longitudinal distance from the odour source. The males of some species of moths have been found to be capable of locating a calling female from a distance of up to 4000m away (Wilson, 1963). While this would be an essential skill were there very few females in a large area, if many females were calling, it would make sense in terms of conservation of energy and time, for a male moth to locate the closest female. Is it possible for a male moth in a pheromone plume to estimate his distance from the calling female by using information from the frequency of interception of odour filaments? The intermittency of a wind tunnel plume (calculated as the proportion of the time when the signal is absent, and which can be translated into an interception frequency) has been shown to decrease as the distance to the source increased (Fackrell and Robins, 1982), suggesting that the frequency of interception of odour filaments would also vary. There is conflicting research, however, carried out in natural odour plumes, that show that the length of odour burst and the time between bursts (which can be converted to intermittency) do not depend greatly on the distance between the plume source and the receiver (Jones, 1983; Murlis and Jones, 1981). It seems that in the more controlled environment of a wind tunnel plume, the frequency at

which odour filaments arrive at a receiver is related to the distance from the source, but in the meandering plumes usually found in natural environments, this frequency information is not a reliable indicator of distance to the odour source.

2. Cross-sectional position in an odour plume. Whilst searching for a calling female, the male moth employs three main, stereotypical search strategies. Firstly, the moth determines the wind direction and flies perpendicular to it in order to make first contact with a pheromone plume. When a pheromone filament is intercepted, the moth performs a "surging" manoeuvre where the main direction of flight is upwind. Should the moth leave the pheromone plume, "casting" behaviour is employed where larger and larger crosswind excursions are performed until either the plume is re-entered, or the moth continues to fly perpendicular to the wind direction, searching for a new pheromone plume (reviewed in Arbas et al., 1993; Cardé and Minks, 1997; Hartlieb and Anderson, 1999 and references therein). After locating a pheromone plume, the optimal strategy for finding the female would be to remain in the plume, and thus reducing the time spent performing "casting" behaviour. Do the odour filaments in the plume provide a mechanism by which this optimal strategy could be performed? In wind tunnel experiments using a stationary EAG, encounters with pheromone filaments in the plume were registered most frequently along the plume centreline, with almost no activity at sampling sites as little as 5 cm from the cross-sectional centre of the plume. The same experimental set-up also found that with an EAG attached to a moth flying in the plume, filaments were encountered most frequently in the centre of the plume, with significantly fewer filament contacts recorded towards the plume edges (Vickers et al., 2001). In a wind tunnel plume, the frequency of interception of odour filaments appears to provide some information on the cross-sectional position within the plume, which could be used by the moth to remain in the plume.

In a natural odour plume, the frequency of arrival of odour pockets is not a reliable measure of the distance from the source of the plume, but the frequency does increase as the centre of the plume is approached (Vickers et al., 2001). Different ORNs capable of following different frequencies of odour interception may help the moth determine where in the cross-sectional area of the plume it currently is, and so remain in the plume for longer periods.

6.5.3.3 Determination of the Blend Components

When searching for a calling female, the male moth is only interested in con-specific females to allow for successful reproduction. The pheromone released by the female is usually a blend of several different components, produced in different ratios. The pheromone blends of some species are very similar, differing only in relative ratios or in a single different component making up the blend (see Tamaki and Yushima, 1974 for an early example of experiments showing different component ratios; and Witzgall et al., 2004 for examples of species with at least one pheromone component in common). So for successful reproduction, the male moth needs to be able to detect the correct pheromone components in the odour plume. These components are detected by different ORNs on the antennae. Although it can be assumed that the female releases all of the pheromone together in pulses, what happens if the constituent components do not arrive simultaneously at the ORNs? Does the moth still detect the correct blend? Taking into account the differences in decay times for ORNs found during this investigation could solve the problem of delays between pheromone component arrivals at the ORNs. The ORNs with longer decays, that continue to fire for longer than the stimulus duration, could create a time window in which other pheromone components could arrive to send information on the blend to the AL and higher brain centres, whilst ORNs with shorter decays could be used to signal to the moth whether it is still in the odour plume or not. While there is currently no evidence that this mechanism of blend detection actually exists in the moth, the efficient detection of the different pheromone compounds at the antennae is important for the moth to allow for the detection of the blend and hence confirm the species of the calling female.

6.6 Conclusions

Electrophysiological studies show that pheromone receptive ORNs have a high selectivity for specific pheromone components, and are a highly sensitive group of neurons. In Chapter 3, this ORN sensitivity is demonstrated, with the convergence of up to ~10000 ORNs at a single site in the AL leading to a further enhancement. This convergence suggests a high level of redundancy in the signals from the ORNs, but there is the possibility that each ORN also provides some information independent of other ORNs, suggesting a reduction in the redundancy in accordance with Barlow's efficient coding hypothesis (Barlow, 1961, 2001). By analysing many repeat recordings from several ORNs in response to the major pheromone component, the work in this chapter suggests that the temporal response of individual ORNs is an independent factor that may be utilised at higher processing centres in the AL. While the frequency analysis of these temporal responses was not sufficient to prove that ORNs are capable of resolving stimuli at specific frequencies, it can be seen that individual ORNs express an interest in different frequency bands, which correspond to the range of frequencies of odour filament interception expected in a plume navigable by a flying moth.

Chapter 7

Conclusions

7.1 Overview of Work Completed

This section provides a brief description of the work carried out in relation to the aims of the thesis, with an assessment of the overall results.

Initial aim: Investigate detection thresholds at the first two stages of the olfactory pathway (ORNs on the antenna and ALNs in the AL) of a moth.

i. A performance measure, the net spikes s^{-1}, was used to determine the numbers of spikes produced by both ORNs and ALNs in response to the stimulus. Distributions of these net spikes s^{-1} were generated for each stimulus load tested, including blank.

ii. Statistical analyses (t-tests and AUC comparisons) were performed to determine which stimulus loads produced distributions that were statistically different to that of the blank. These statistical analyses provided estimates of the detection thresholds for the ORNs and ALNs. The difference between these thresholds gave an indication of the boost in sensitivity achieved.

iii. A theoretical model assuming Poisson ORNs and a simple summation at a convergence site (PN) suggested that the sensitivity enhancement in the biological system exceeded that explained by the model.

iv. Several hypotheses to explain the observed boost were suggested:

a) Temporal Encoding: The encoding method of the ORNs may not be rate based at all, but could rely on patterns of spikes.

b) Active Convergence: The glomerulus may actually be an active site of ORN convergence, capable of boosting the signal in response to the stimulus.

c) Lucky ORN: Due to the quantal nature of odour stimuli, not all ORNs may receive enough stimulus molecules to initiate a response, while others do.

d) Synchronous Spikes: Individual ORNs may be capable of synchronising their responses to the stimulus, and so would require only a single spike each to signal the presence of a stimulus.

This aim was largely completed, with cogent statistical analyses used to assess the detection thresholds, and thus quantify the transmission of olfactory information from periphery to a more central structure of the nervous system. The addition of a theoretical model allowed for a direct comparison of analytical results with empirical electrophysiological data. In this case, the result was greater than the sum of the parts, suggesting that more complex coding strategies are at work during early olfaction in this animal.

Further aim: Investigate ways in which this boost could be achieved. For the first part of this aim, the encoding scheme at the periphery, hypothesis a), was investigated.

v. If the presence of a stimulus could be coded for by a pattern of spikes, the number of spikes necessary could be too few to be detected by the rate metric that was employed to calculate the detection thresholds. A preliminary investigation into the type of encoding scheme used by the ORNs could not discount the possibility of a temporal code.

vi. An investigation into the temporal structures of spike patterns across repeat recordings from individual neurons was performed using UEA. Unfortunately, results were inconclusive due to the strong non-stationary nature of these spike trains in response to a stimulus presentation.

vii. An investigation into the spike patterns across a population of neurons was also inconclusive, due to the incompatibility of the data with the statistical basis of underlying the method employed (UEMWA).

An investigation into the encoding scheme at the periphery was performed. Several methods of assessing this scheme were employed, but were not capable of providing conclusive evidence for either a rate or temporal encoding scheme. While the results from this section were inconclusive, it does show that care must be taken when selecting methods for the investigation of encoding schemes since some methods assume the data adheres to specific structural formats, and so may not be well suited to all data sets.

For the second part of this aim, a model of a possible glomerular convergence site, related to hypothesis b) was created and used to investigate whether the detection of the other two hypotheses, c) lucky ORN and d) synchronous spikes, could be enhanced.

viii. A computational model was created using real ORN input spike statistics and an integrate-and-fire output PN. Simulations showed that this more realistic model could not account for the detection ability of the biological system.

ix. If a glomerulus were to act as an active site of convergence of ORNs, a multitude of different calculations could be performed. Here, a model was created incorporating subunits representing convergence sites of subsets of ORNs at separate dendritic segments of a PN. Each subunit performed a nonlinear thresholding function on the input spike trains, passing the output onto an integrate-and-fire PN.

x. To assess the nonlinear model, a comparison model was created incorporating all the features of the nonlinear model, but with each subunit performing a linear summation of the input spikes.

xi. To investigate hypothesis c), subsets of input ORNs were designated "lucky" i.e. only this subset showed a response to a stimulus, in terms of an increase in firing rate, while the remainder of the population did not. In this case, the detection performance of the nonlinear model exceeded that of the linear version.

xii. To investigate hypothesis d), subsets of input ORNs again responded to a stimulus, but with a single spike each, time-locked with those of other responding ORNs. The nonlinear model far outperformed the linear model in terms of detecting the stimulus, even with the addition of timing jitter on these synchronous spikes.

A nonlinear model was developed that outperformed a linear model, in terms of stimulus detection ability of the output neuron, in two specific areas, when only a subset of input ORNs responded to a stimulus presentation, and when ORNs responded with a single synchronous spike. While the actual model incorporated features from different types of neurons found in different animals, it still suggests that site of receptor convergence in the moth olfactory pathway is likely to be active (i.e. capable of transforming input signals in some way), rather than a simple, passive summation of the inputs.

As an additional section, an initial investigation of the frequency dynamics of receptor neurons was performed to assess the possibility that ORNs were capable of coding for different frequencies of stimulus interception, and so providing an independent factor reducing the redundancy in the convergence of many ORNs. While this work assumed a simple first-order linear description of an ORN, and used a single square pulse stimulus, different frequency dynamics were observed, suggesting that further investigation might provide useful insights into the coding capacity of these neurons.

7.2 Further Areas of Work

Temporal Coding: Only a couple of methods were implemented here for the investigation of temporal patterns in spike trains. Although it is quite likely that the encoding

scheme at the periphery is rate based, there are various other methods currently available which may find a temporal scheme. Another interesting area would be to take recordings simultaneously from several different ORNs.

Modelling: On the computational side, further experimentation with the shape of thresholding function within the subunits may provide further enhancement of the detection capability of the output neuron. Given the multitude of synaptic connections within each glomerulus, a myriad of other structural arrangements may provide insight into the boost observed in the biology. On the biological side, it would be interesting to further investigate the synaptic connectivity within a moth glomerulus, using a variety of staining techniques, in order to produce a more biologically realistic computational model.

Frequency Analysis: The use of a white noise, or variable sequence stimulus would enable further analysis of the frequency encoding capabilities of ORNs.

7.3 Outlook

Due to the accessibility and genetic manipulations possible in this system, olfaction is an ideal pathway in which to study the efficient transmission of neural information, particularly where converging architecture is involved. But in order to understand the brain and its coding mechanisms, an integrative approach is necessary utilising a combination of both electrophysiological and computational methodologies.

Bibliography

M. Abeles, H. Bergman, E. Margalit, and E. Vaadia. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, 70:1629–1638, 1993.

A.M.H.J. Aertsen, G.L. Gerstein, M.K. Habib, and G. Palm. Dynamics of neuronal firing correlation: Modulation of "effective connectivity". *Journal of Neurophysiology*, 61:900–917, 1989.

H. Ai and R. Kanzaki. Modular organization of the silkmoth antennal lobe macroglomerular complex revealed by voltage-sensitive dye imaging. *Journal of Experimental Biology*, 207: 633–644, 2004.

A.C. Allison and R.T.T. Warwick. Quantitative observations on the olfactory system of the rabbit. *Brain*, 72:186–197, 1949.

P. Anderson, M. Hilker, B.S. Hansson, S. Bombosch, B. Klein, and H. Schildknecht. Oviposition deterring components in larval frass of spodoptera littoralis (boisd.) (lepidoptera: Noctuidae): A behavioural and electrophysiological evaluation. *Journal of Insect Physiology*, 39:129–137, 1993.

P. Anderson, B.S. Hansson, and J. Löfqvist. Plant-odour-specific receptor neurons on the antennae of female and male spodoptera littoralis. *Physiological Entomology*, 20:189–198, 1998.

A.M. Angioy, A. Desogus, I.T. Barbarossa, P. Anderson, and B.S. Hansson. Extreme sensitivity in an olfactory system. *Chemical Senses*, 28:279–284, 2003.

S. Anton and B.S. Hansson. Central processing of sex pheromone, host odour, and oviposition deterrent information by interneurons in the antennal lobe of female spodoptera littoralis (lepidoptera: Noctuidae). *Journal of Comparative Neurology*, 350:199–214, 1994.

S. Anton and B.S. Hansson. Sex pheromone and plant-associated odour processing in antennal lobe interneurons of male spodoptera littoralis (lepidoptera: Noctuidae). *Journal of Comparative Physiology A*, 176:773–789, 1995.

S. Anton and U. Homberg. Antennal lobe structure. In B.S. Hansson, editor, *Insect Olfaction*, pages 97–124. Springer-Verlag, 1999.

E.A. Arbas, M.A. Willis, and R. Kanzaki. Organization of goal-oriented locomotion: pheromone-modulated flight behavior of moths. In R.D. Beer, R.E. Ritzmann, and T. McKenna, editors, *Biological Neural Networks in Invertebrate Neuroethology and Robotics*, pages 159–198. Academic Press, San Diego, 1993.

K.A. Archie and B.W. Mel. An intradendritic model for computation of binocular disparity. *Nature Neuroscience*, 3:54–63, 2000.

T.C. Baker and K.F. Haynes. Field and laboratory electroantennographic measurements of pheromone plume structure correlated with oriental fruit moth behavior. *Physiological Entomology*, 14:1–12, 1989.

T.C. Baker, M.A. Willis, K.F. Haynes, and P.L. Phelan. A pulsed cloud of sex pheromone elicits upwind flight in male moths. *Physiological Entomology*, 10:257–265, 1985.

D. Bamber. The area above the ordinal dominance graph and the area below the receiver operating graph. *Journal of Mathematical Psychology*, 12:387–415, 1975.

H.B. Barlow. Possible principles underlying the transformation of sensory messages. In W.A. Rosenblith, editor, *Sensory communication*, pages 217–234. MIT Press, Cambridge, MA, 1961.

H.B. Barlow. Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12:241–253, 2001.

H.B. Barlow, W.R. Levick, and M. Yoon. Responses to single quanta of light in retinal ganglion cells. *Vision Res. Suppl.*, 3:87–101, 1971.

J. Bau, K.A. Justus, and R.T. Cardé. Antennal resolution of pulsed pheromone plumes in three moth species. *Journal of Insect Physiology*, 48:433–442, 2002.

D.A. Baylor, B.J. Nunn, and J.L. Schnapf. The photocurrent, noise and spectral sensitivity of rods of the monkey macaca fascicularis. *Journal of Physiology*, 357:575–607, 1984.

B. G. Berg, T. J. Almaas, J. G. Bjaalie, and H. Mustaparta. The macroglomerular complex of the antennal lobe in the tobacco budworm moth heliothis virescens: specified subdivision in four compartments according to information about biologically significant compounds. *Journal of Comparative Physiology A*, 183:669–682, 1998.

J.C. Bezdek. *Pattern recognition with fuzzy objective function algorithms.* New York: Plenum, 1981.

K.P. Bhatnager and C.K. Kallen. Quantitative observations on the nasal epithelia and olfactory innervation in bats. *Acta Anatomy*, 91:272–282, 1975.

J. Boeckh and K.D. Ernst. Contribution of single unit analysis in insects to an understanding of olfactory function. *Journal of Comparative Physiology A*, 161:549–565, 1987.

J. Boeckh, K.D. Ernst, H. Saas, and U. Waldow. Anatomical and physiological characteristics of individual neurones in the central antennal pathway of insects. *Journal of Insect Physiology*, 30:15–26, 1984.

L.B. Buck. Information coding in the vertebrate olfactory system. *Annual Review of Neuroscience*, 19:517–544, 1996.

D.V. Buonomano and M.M. Merzenich. Temporal information transformed into a spatial code by a network with realistic properties. *Science*, 267:1028–1030, 1995.

W. Calvin and C. Stevens. Synaptic noise and other sources of randomness in motoneuron interspike intervals. *Journal of Neurophysiology*, 31:574–587, 1968.

R.T. Cardé and A.K. Minks. *Insect Pheromone Research: New Directions.* Chapman and Hall, New York., 1997.

R.T. Cardé and W.L. Roelofs. Temperature modification of male sex pheromone response and factors affecting female calling in holomelina immaculate (lepidoptera: Arctiidae). *Canadian Entomologist*, 105:1505–1512, 1973.

R.T. Cardé, L.L. Dindonis, B. Agar, and J. Foss. Apparency of pulsed and continuous pheromone to male gypsy moths. *Journal of Chemical Ecology*, 10:335–348, 1984.

M.A. Carlsson and B.S. Hansson. Responses in highly selective sensory neurons to blends of pheromone components in the moth agrotis segetum. *Journal of Insect Physiology*, 48: 443–451, 2002.

M.A. Carlsson, C.G. Galizia, and B.S. Hansson. Spatial representation of odours in the antennal lobe of the moth spodoptera littoralis (lepidoptera: Noctuidae). *Chemical Senses*, 27:231–244, 2002.

W.R. Chen and G.M. Shepherd. Putting odor maps in sync. *Nature Neuroscience*, 5: 505–506, 2002.

A. Chess, I. Simon, H. Cedar, and R. Axel. Allelic inactivation regulates olfactory receptor gene expression. *Cell*, 78:823–834, 1994.

K.Y. Chong. Antennal lobe modelling of the moth spodoptera littoralis. First Year PhD Report. Unpublished, 2003.

T.A. Christensen and J.G. Hildebrand. Male-specific, sex pheromone-selective projection neurons in the antennal lobes of the moth manduca sexta. *Journal of Comparative Physiology A*, 160:553–569, 1987.

T.A. Christensen and J.G. Hildebrand. Frequency coding by central olfactory neurons in the sphinx moth manduca sexta. *Chemical Senses*, 13:123–130, 1988.

T.A. Christensen and J.G. Hildebrand. Coincident stimulation with pheromone components improves temporal pattern resolution in central olfactory neurons. *Journal of Neurophysiology*, 77:775–781, 1997.

T.A. Christensen and J.G. Hildebrand. Pheromonal and host-odor processing in the insect antennal lobe: How different? *Current Opinions in Neuroscience*, 12:393–399, 2002.

T.A. Christensen, H. Mustaparta, and J.G. Hildebrand. Discrimination of sex pheromone blends in the olfactory system of the moth. *Chemical Senses*, 14:463–477, 1989.

T.A. Christensen, B.R. Waldrop, I.D. Harrow, and J.G. Hildebrand. Local interneurons and information processing in the olfactory glomeruli of the moth manduca sexta. *Journal of Comparative Physiology A*, 173:385–399, 1993.

T.A. Christensen, T. Heinbockel, and J.G. Hildebrand. Olfactory information processing in the brain: encoding chemical and temporal features of odors. *Journal of Neurobiology*, 30:82–91, 1996.

T.A. Christensen, V.M. Pawlowski, H. Lei, and J.G. Hildebrand. Multi-unit recordings reveal context-dependent modulation of synchrony in odor-specific neural ensembles. *Nature Neuroscience*, 3:927–931, 2000.

T.A. Christensen, G. D'Alessandro, J. Lega, and J.G. Hildebrand. Morphometric modeling of olfactory circuits in the antennal lobe: I. simulations of spiking local interneurons. *Biosystems*, 61:143–153, 2001.

W.E. Conner, T. Eisner, K. Vander Meer, R, A. Guerrero, D. Ghiringelli, and J. Meinwald. Sex attractant of an arctiid moth (utethesia ornatrix): a pulsed chemical signal. *Behavioral Ecology and Sociobiology*, 7:55–63, 1980.

J.W. Cooley and J.W. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19:297–301, 1965.

W.W. Cooley and P.R. Lohnes. *Multivariate data analysis.* Wiley, New York, 1971.

C.T. David, J.S. Kennedy, and A.R. Ludlow. Finding of a sex pheromone source by gypsy moths released in the field. *Nature*, 303:804–806, 1983.

P. Dayan and L.F. Abbott. *Theoretical neuroscience: computational and mathematical modeling of neural systems.* MIT Press, Cambridge, USA, 2001.

A. Destexhe, Z.F. Mainen, and T.J. Sejnowski. Synaptic currents, neuromodulation and kinetic models. In M.A. Arbib, editor, *The handbook of brain theory and neural networks.*, pages 956–959. MIT Press, Cambridge, USA, 1998.

M. Diesmann, M.-O. Gewaltig, and A. Aertsen. Conditions for stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402:529–533, 1999.

P.G. Distler. Histochemical demonstration of gaba-like immunoreactivity in colbalt labeled neuron individuals in the insect olfactory pathway. *Histchemistry*, 91:245–249, 1989.

P.G. Distler and J. Boeckh. Synaptic connections between olfactory receptor cells and uniglomerular projection neurons in the antennal lobe of the american cockroach, periplaneta americana. *Journal of Computational Neurology*, 370:35–46, 1996.

P.G. Distler and J. Boeckh. Synaptic connections between identified neuron types in the antennal lobe glomeruli of the cockroach, periplaneta americana: I. uniglomerular projection neurons. *Journal of Comparative Neurology*, 378:307–319, 1997a.

P.G. Distler and J. Boeckh. Synaptic connections between identified neuron types in the antennal lobe glomeruli of the cockroach, periplaneta americana: Ii. local multiglomerular interneurons. *Journal of Comparative Neurology*, 383:529–540, 1997b.

P.G. Distler, C. Gruber, and J. Boeckh. Synaptic connections between gaba-immunoreactive neurons and uniglomerular projection neurons within the antennal lobe of the cockroach, periplaneta americana. *Synapse*, 29:1–13, 1998.

N.R. Draper and H. Smith. *Applied regression analysis, 3rd edition.* John Wiley and Sons, New York., 1998.

P. Duchamp-Viret, A. Duchamp, and M. Vigouroux. Amplifying role of convergence in olfactory system. a comparative study of receptor cell and second-order neuron sensitivities. *Journal of Neurophysiology*, 61:1085–1094, 1989.

D. Dumitrescu, B. Lazzareni, and L.C. Jain. *Fuzzy sets and their application to clustering and training.* New York: CRC, 2000.

S. Everitt and G. Dunn. *Applied multivariate data analysis.* Edward Arnold, London, 1999.

J.E. Fackrell and A.G. Robins. The effect of source size on concentration fluctuations in plumes. *Boundary-Layer Meteorology*, 22:335–350, 1982.

J.-M. Fellous, P.H.E. Tiesinga, P.J. Thomas, and T.J. Sejnowski. Discovering spike patterns in neuronal responses. *Journal of Neuroscience*, 24:2989–3001, 2004.

G.D. Field and F. Rieke. Nonlinear signal transfer from mouse rods to bipolar cells and implications for visual sensitivity. *Neuron*, 34:773–785, 2002.

G.D. Field, A.P. Sampath, and F. Rieke. Retinal processing near absolute threshold: From behavior to mechanism. *Annual Review of Physiology*, 67:491–514, 2005.

F. Gabbiani and C. Koch. Principles of spike train analysis. In C. Koch and I. Segev, editors, *Methods in neuronal modeling: from ions to networks.*, pages 313–360. MIT Press, Cambridge, USA, 1998.

A. Gelperin. Oscillatory dynamics and information processing in olfactory systems. *Journal of Experimental Biology*, 202:1855–1864, 1999.

G. Gemne and K.B. Døving. Ultrastructural properties of olfactory neurons in fish (lota lota l.). *American Journal of Anatomy*, 126:457–463, 1969.

G. Gerstein and B. Mandelbrot. Random walk models for the spike activity of a single neuron. *Journal of Biophysics*, 4:41–68, 1964.

C.M. Gray, P. Konig, A.K. Engel, and W. Singer. Oscillatory responses in cat visual-cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337, 1989.

D. Green and J. Swets. *Signal detection theory and psychophysics.* John Wiley and Sons, New York, 1966.

S. Grün. *Unitary joint-events in multiple-neuron spiking activity: Detection, significance and interpretation.* Reihe Physik, Band 60. Then, Frankfurt/Main: Verlag Harri Deutsch., 1996.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: I. detection and significance. *Neural Computation*, 14:43–80, 2001a.

S. Grün, M. Diesmann, and A. Aertsen. Unitary events in multiple single-neuron spiking activity: Ii. nonstationary data. *Neural Computation*, 14:81–119, 2001b.

A.L. Gustavsson, M. Tuvesson, M.C. Larsson, W. Wenqi, B.S. Hansson, and T. Liljefors. Bioisosteric approach to elucidation of binding of the acetate group of a moth sex pheromone component to its receptor. *Journal of Chemical Ecology*, 23:2755–2776, 1997.

J.A. Hanley and B.J. McNeil. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143:29–36, 1982.

J.A. Hanley and B.J. McNeil. A method of comparing the areas under receiver operating characteristic curves derived from the same cases. *Radiology*, 148:839–843, 1983.

B. S. Hansson, T. J. Almaas, and S. Anton. Chemical communication in heliothine moths. 5. antennal lobe projection patterns of pheromone-detecting olfactory receptor neurons in the male heliothis virescens (lepidoptera: Noctuidae). *Journal of Comparative Physiology A*, 177:535–543, 1995.

B.S. Hansson. Olfaction in Lepidoptera. *Experientia*, 51:1003–1027, 1995.

B.S. Hansson. A bug's smell - research into insect olfaction. *Trends in Neurosciences*, 25: 270–274, 2002.

B.S. Hansson and S. Anton. Function and morphology of the antennal lobe: New developments. *Annual Review of Entomology*, 45:203–231, 2000.

B.S. Hansson and T.A. Christensen. Functional characteristics of the antennal lobe. In B.S. Hansson, editor, *Insect Olfaction*, pages 126–161. Springer-Verlag, 1999.

B.S. Hansson, J.N.C. Van der Pers, and J. Löfqvist. Comparison of male and female olfactory cell response to pheromone compounds and plant volatiles in the turnip moth, agrotis segetum. *Physiological Entomology*, 14:147–155, 1989.

B.S. Hansson, H. Ljungberg, E. Hallberg, and C. Löfstedt. Functional specialization of olfactory glomeruli in a moth. *Science*, 256:1313–1315, 1992.

F.J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66:51–84, 1978.

E. Hartlieb and P. Anderson. Olfactory-released behaviours. In B.S. Hansson, editor, *Insect olfaction.*, pages 315–349. Springer-Verlag, 1999.

E. Hartlieb, S. Anton, and B.S. Hansson. Dose-dependent response characteristics of antennal lobe neurons in the male moth agrotis segetum (lepidoptera: Noctuidae). *Journal of Comparative Physiology A*, 181:469–476, 1997.

M. Häusser, N. Spruston, and G.J. Stuart. Diversity and dynamics of dendritic signalling. *Science*, 290:739–744, 2000.

J.G. Hildebrand. Olfactory control of behavior in moths: Central processing of odor information and the functional significance of olfactory glomeruli. *Journal of Comparative Physiology A*, 178:5–19, 1996.

C.F. Hinks and J.R. Byers. Biosystematics of the genus euxoa (lepidoptera: Noctuidae). v. rearing procedures and life cycles of 36 species. *Canadian Entomologist*, 108:1345–1357, 1976.

U. Homberg, T.A. Christensen, and J.G. Hildebrand. Structure and function of the deutocerebrum in insects. *Annual Review of Entomology*, 34:477–501, 1989.

F. Hoppner, R. Kruse, F. Klawonn, and T. Runkler. *Fuzzy cluster analysis: methods for classification, data analysis and image recognition.* New York: Wiley, 1999.

M. Hösl. Pheromone-sensitive neurons in the deutocerebrum of periplaneta americana: receptive fields on the antenna. *Journal of Comparative Pnysiology A.*, 167:321–327, 1990.

J.C. Hsu. *Multiple Comparisons: Theory and methods.* Chapman and Hall, London, 1996.

D.H. Hubel. Tungsten microelectrode for recording from single units. *Science*, 125:549–550, 1957.

J.E. Jackson. *A user's guide to principal components.* Wiley, New York, 1988.

C.D. Jones. On the structure of instantaneous plumes in the atmosphere. *Journal of Hazardous Materials*, 7:87–112, 1983.

M. Jönsson and P. Anderson. Electrophysiological response to herbivore induced host plant volatiles in the moth spodoptera littoralis. *Physiological Entomology*, 24:377–385, 1999.

K.A. Justus and R.T. Cardé. Flight behaviour of males of two moths, cadra cautella and pectinophora gossypiella, in homogeneous clouds of pheromone. *Physiological Entomology*, 27:67–75, 2002.

K.A. Justus, J. Murlis, C. Jones, and R.T. Cardé. Measurement of odor-plume structure in a wind tunnel using a photoionization detector and a tracer gas. *Environmental Fluid Mechanics*, 2:115–142, 2002a.

K.A. Justus, S.W. Schofield, J. Murlis, and R.T. Cardé. Flight behavior of cadra cautella males in rapidly pulsed pheromone plumes. *Physiological Entomology*, 27:58–66, 2002b.

K.A. Justus, R.T. Cardé, and A.S. French. Dynamic properties of antennal responses to pheromone in two moth species. *Journal of Neurophysiology*, 93:2233–2239, 2005.

K.-E. Kaissling. Insect olfaction. In L.M. Beidler, editor, *Handbook of Sensory Physiology*, pages 351–431. Springer-Verlag, 1971.

K.-E. Kaissling. Temporal characteristics of pheromone receptor cell responses in relation to orientation behavour of moths. In T.L. Payne, M.C. Birch, and C.E.J. Kenney, editors, *Mechanisms in insect olfaction.*, pages 5–47. Clarendon Press, Oxford, 1986a.

K.-E. Kaissling and E. Priesner. Die riechschwelle des seidenspinners (the olfactory threshold of silkworm moths). *Die Naturwissenschaften*, 57:23–28, 1970.

K.E. Kaissling. Chemo-electrical transduction in insect olfactory receptors. *Annual Review of Neuroscience*, 9:121–145, 1986b.

K.E. Kaissling, J.G. Hildebrand, and J.H. Tumlinson. Pheromone receptor cells in the male moth manduca sexta. *Archives of Insect Biochemical Physiology*, 10:273–264, 1989.

B. Kalinovà, M. Hoskovec, I. Liblikas, C.R. Unelius, and B.S. Hansson. Detection of sex pheromone components in manduca sexta (l.). *Chemical Senses*, 26:1175–1186, 2001.

R. Kanzaki, E.A. Arbas, and N.J. Strausfeld. Physiology and morphology of projection neurons in the antennal lobe of the male moth manduca sexta. *Journal of Comparative Physiology A*, 165:427–453, 1989.

G. Kasang. Physiochemical events in olfaction of the silkmoth. *Naturwissenshaften*, 60: 95–101, 1973.

S.M. Kay and S.J.Jr. Marple. Spectrum analysis - a modern perspective. *Proceedings of the IEEE*, 69:1380–1419, 1981.

T.A. Keil. Morphology and development of the peripheral olfactory organs. In B.S. Hansson, editor, *Insect Olfaction*, pages 5–47. Springer-Verlag, 1999.

T.A. Keil and R.A. Steinbrecht. Diffusion barriers in silkmoth sensory epithelia: application of lanthanum tracer to olfactory sensilla of antheraea polyphemus and bombyx mori. *Tissue Cell*, 19:119–134, 1987.

J.R. King, T.A. Christensen, and J.G. Hildebrand. Response characteristics of an identified, sexually dimorphic olfactory glomerulus. *Journal of Neuroscience*, 20:2391–2399, 2000.

S.A. Klein and D.M. Levi. Hyperacuity thresholds of 1 sec: theoretical predictions and empirical validation. *Journal of the Optical Society of America*, 2:1170–1190, 1985.

C. Koch. *Biophysics of Computation: Information Processing in Single Neurons.* Oxford University Press, Oxford, 1999.

C. Koch, T. Poggio, and V. Torre. Retinal ganglion cells: A functional interpretation of dendritic morphology. *Philosophical Transactions of the Royal Society of London. B Biological Sciences*, 298:227–264, 1982.

B. Kodadová. Resolution of pheromone pulses in receptor cells of antheraea polyphemus at different temperatures. *Journal of Comparative Physiology A*, 179:301–310, 1996.

M.A. Koontz and D. Schneider. Sexual dimorphism in neuronal projections from the antennae of silk moths (bombyx mori, antheraea polyphemus) and the gypsy moth (lymantria dispar). *Cell and Tissue Research*, 249:39–50, 1987.

W.J. Krzanowski. *Principles of multivariate analysis.* Clarendon Press, Oxford, 1988.

S.B. Laughlin. A simple coding procedure enhances a neuron's information capacity. *Zeitschrift fr Naturforschung C*, 36:910–912, 1981.

G. Laurent. A systems perspective on early olfactory coding. *Science*, 286:723–728, 1999.

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

G. Laurent, M. Wehr, and H. Davidowitz. Temporal representations of odors in an olfactory network. *Journal of Neuroscience*, 16:3837–3847, 1996.

G. Laurent, M. Stopfer, R.W. Friedrich, M.I. Rabinovich, A. Volkovskii, and H.D.I. Abarbanel. Odor coding as an active, dynamical process: Experiments, computation and theory. *Annual Review of Neuroscience*, 24:263–297, 2001.

H. Lei, T.A. Christensen, and J.G. Hildebrand. Local inhibition modulates odor-evoked synchronization of glomerulus-specific output neurons. *Nature Neuroscience*, 5:557–565, 2002.

H. Lei, T.A. Christensen, and J.G. Hildebrand. Spatial and temporal organization of ensemble representations for different odor classes in the moth antennal lobe. *Journal of Neuroscience*, 24:11108–11119, 2004.

W.C. Lemon and W.M. Getz. Rate code input produces temporal code output from cockroach antennal lobes. *BioSystems*, 58:151–158, 2000.

R.L. Lerner, T.K. Gyorgyi, J. Reagan, A. Roby-Shemkovitz, R. Rybeynski, and R.G. Vogt. Peripheral events in moth olfaction. *Chemical Senses*, 15:191–198, 1990.

M.S. Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network - Computation in Neural Systems*, 9:R53–R78, 1998.

H. Ljungberg, P. Anderson, and B.S. Hansson. Physiology and morphology of pheromone-specific sensilla on the antennae of male and female spodoptera littoralis (lepidoptera: Noctuidae). *Journal of Insect Physiology*, 39:253–260, 1993.

R. Llinás and C. Nicholson. Electrophysiological properties of dendrites and somata in alligator purkinje cells. *Journal of Neurophysiology*, 34:534–551, 1971.

K. MacLeod, A. Bäcker, and G. Laurent. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature*, 395:693–698, 1998.

A. Mafra-Neto and R.T. Cardé. Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Nature*, 369:142–144, 1994.

A. Mafra-Neto and R.T. Cardé. Influence of plume structure and pheromone concentration on upwind flight of cadra cautella males. *Physiological Entomology*, 20:117–133, 1995a.

A. Mafra-Neto and R.T. Cardé. Effect of the fine-scale structure of pheromone plumes: pulse frequency modulates activation and upwind flight of almond moth males. *Physiological Entomology*, 20:229–242, 1995b.

Z.F. Mainen and T.J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503–1506, 1995.

R.W. Mankin, M.S. Mayer, and A.J. Grant. Sensitivity of trichoplusia ni (hübner) pheromone receptor neurons: Relationships between neural thresholds and behavioural responses. *Journal of Comparative Physiology A*, 168:739–747, 1991.

D.J. Mar, C.C. Chow, W. Gerstner, R.W. Adams, and J.J. Collins. Noise shaping in populations of coupled model neurons. *Proceedings of the National Academy of Science, USA*, 96:10450–10455, 1999.

F. Marion-Poll and T.R. Tobin. Temporal coding of pheromone pulses and trains in manduca sexta. *Journal of Comparative Physiology A*, 171:505–512, 1992.

M.V. Mascagni and A.S. Sherman. Numerical methods for neuronal modeling. In C. Koch and I. Segev, editors, *Methods in neuronal modeling: from ions to networks.*, pages 568–604. MIT Press, Cambridge, USA, 1998.

D.N. Mastronarde. Correlated firing of cat retinal ganglion cells. ii responses of x- and y-cells to single quantal events. *Journal of Neurophysiology*, 49:325–349, 1983.

J. Maynard-Smith. *Evolutionary genetics.* Oxford University Press, Oxford, 1998.

W. McCullough and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5:115–133, 1943.

B.W. Mel. The clusteron: Toward a simple abstraction for a complex neuron. In J. Moody, S. Hanson, and R. Lippmann, editors, *Advances in neural information processing systems, volume 4*, pages 35–42. Morgan Kaufmann, San Mateo, California, 1992a.

B.W. Mel. Nmda-based pattern discrimination in a modeled cortical neuron. *Neural Computation*, 4:502–516, 1992b.

B.W. Mel. Synaptic integration in an excitable dendritic tree. *Journal of Neurophysiology*, 70:1086–1101, 1993.

B.W. Mel, D.I. Ruderman, and K.A. Archie. Translation-invariant orientation tuning in visual 'complex' cells could derive from intradendritic computation. *Journal of Neuroscience*, 18:4325–4334, 1998.

A.V. Minor and K.-E. Kaissling. Cell responses to single pheromone molecules may reflect the activation kinetics of olfactory receptor molecules. *Journal of Comparative Physiology A*, 189:221–230, 2003.

P.A. Moore and J. Atema. Spatial information in the three-dimensional fine structure of an aquatic odor plume. *Biological Bulletin*, 181:408–418, 1991.

J. Murlis. The structure of odour plumes. In T.L. Payne, M.C. Birch, and C.E.J. Kennedy, editors, *Mechanisms in insect olfaction.*, pages 351–431. Oxford University Press, New York., 1986.

J. Murlis and C.D. Jones. Fine-scale structure of odour plumes in relation to insect orientation to distant pheromone and other attractant sources. *Physiological Entomology*, 6: 71–86, 1981.

J. Murlis, J.S. Elkington, and R.T. Cardé. Odor plumes and how insects use them. *Annual Review of Entomology*, 37:505–532, 1992.

J. Murlis, M.A. Willis, and R.T. Cardé. Spatial and temporal structures of pheromone plumes in fields and forests. *Physiological Entomology*, 25:211–222, 2000.

H. Mustaparta. Central mechanisms of pheromone information processing. *Chemical Senses*, 21:269–275, 1996.

J. Ngai, A. Chess, M.M. Dowling, N. Necles, E.R. Macagno, and R. Axel. Coding of olfactory information: Topography of odorant receptor expression in the catfish olfactory epithelium. *Cell*, 72:667–680, 1993.

S. Nirenberg, S.M. Carcieri, A.L. Jacobs, and P.E. Latham. Retinal ganglion cells act largely as independent encoders. *Nature*, 411:698–701, 2001.

H. Nyquist. Certain topics in telegraph transmission theory. In *Transactions of the American Institute of Electrical Engineers*, pages 617–644, 1928.

S. A. Ochieng', P. Anderson, and B. S. Hansson. Al projection patterns of olfactory receptor neurons involved in sex pheromone detection in spodoptera littoralis (lepidoptera: Noctuidae). *Tissue Cell*, 27:221–232, 1995.

T.C. Pearce, P.F.M.J. Verschure, J. White, and J.S. Kauer. Robust stimulus encoding in olfactory processing: hyperacuity and efficient signal transmission. In S. Wermter, J. Austin, and D. Willshaw, editors, *Emergent neural computation architectures based on neuroscience.*, pages 461–479. Springer-Verlag, 2001a.

T.C. Pearce, P.F.M.J. Verschure, J. White, and J.S. Kauer. Stimulus encoding during the early stages of olfactory processing: a modeling study using an artificial olfactory system. *Neurocomputing*, 38:399–406, 2001b.

P. Poirazi, T.M. Brannon, and B.W. Mel. Arithmetic of sub-threshold synaptic summation in a model of ca1 pyramidal cell. *Neuron*, 37:977–987, 2003a.

P. Poirazi, T.M. Brannon, and B.W. Mel. Pyramidal neuron as two-layer neural network. *Neuron*, 37:989–999, 2003b.

W. Rall and I. Segev. Functional possibilities for synapses on dendrites and on dendritic spines. In G. Edelman, W. Gall, and W. Cowan, editors, *Synaptic function*, pages 605–636. Wiley, New York, 1987.

A. Riehle, S. Grün, M. Diesmann, and A. Aersen. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278:1950–1953, 1997.

F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the neural code.* MIT Press, Cambridge, Massachusetts., 1997.

F. Rosenblatt. *Principles of neurodynamics.* Spartan, New York, 1962.

J.P. Rospars. Invariance and sex-specific variations of the glomerular organization in the antennal lobes of a moth, mamestra brassicae, and a butterfly, pieris brassicae. *Journal of Comparative Neurology*, 220:80–94, 1983.

J.P. Rospars and J.G. Hildebrand. Anatomical identification of glomeruli in the antennal lobes of the male sphinx moth manduca sexta. *Cell Tissue Research*, 270:205–227, 1992.

A. Roy, P. Steinmetz, and E. Niebur. Rate limitations of unitary event analysis. *Neural Computation*, 12:2063–2082, 2000.

E.R. Rumbo and K.E. Kaissling. Temporal resolution of odour pulses by three types of pheromone receptor cell in antheraea polyphemus. *Journal of Comparative Physiology A*, 165:281–291, 1989.

D. Rumelhart, G. Hinton, and J. McClelland. A general framework for parallel distributed processing. In D. Rumelhart and J. McClelland, editors, *Parallel distributed processing: Explorations in the microstructure of cognition, Volume 1*, pages 45–76. Bradford, Cambridge, Massachusetts, 1986.

S. Sachse and C.G. Galizia. Role of inhibition for temporal and spatial odor representation in olfactory output neurons: A calcium imaging study. *Journal of Neurophysiology*, 87: 1106–1117, 2002.

M.M. Sadek, B.S. Hansson, J.P. Rospars, and S. Anton. Glomerular representation of plant volatiles and sex pheromone components in the antennal lobe of the female spodoptera littoralis. *Journal of Experimental Biology*, 205:1636–1376, 2002.

C. Schal and R.T. Cardé. Rhythmic extrusion of pheromone gland elevates pheromone release rate. *Experientia*, 41:1617–1619, 1985.

J. Schiller, G. Major, H.J. Koester, and Y. Schiller. Nmda spikes in basal dendrites of cortical pyramidal neurons. *Nature*, 404:285–289, 2000.

A.M. Schneiderman, J.G. Hildebrand, M.M. Brennan, and J.H. Tumlinson. Trans-sexually grafted antennae alter pheromone-directed behaviour in a moth. *Nature*, 323:801–803, 1986.

G.A.F. Seber. *Multivariate observations.* Wiley, New York, 1984.

I. Segev and M. London. Untangling dendrites with quantitative models. *Science*, 290: 744–750, 2000.

M.N. Shadlen and W.T. Newsome. Noise, neural codes and cortical organization. *Current Opinions in Neurobiology*, 4:569–579, 1994.

M.N. Shadlen and W.T. Newsome. The variable discharge of cortical neurons: Implications for connectivity, computation and information coding. *Journal of Neuroscience*, 18:3870–3896, 1998.

C.E. Shannon. Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37:10–21, 1949.

G. Shepherd and R. Brayton. Logic operations are properties of computer-simulated interactions between excitable dendritic spikes. *Neuroscience*, 21:151–166, 1987.

G.M. Shepherd and C. Koch. Introduction to synaptic circuits. In G.M. Shepherd, editor, *The synaptic organisation of the brain.*, pages 1–36. Oxford University Press, New York, 1998.

R.A. Silver, S.F. Traynelis, and S.G. Cull-Candy. Rapid-time-course miniature and evoked excitatory currents at cerebellar synapses in situ. *Nature*, 355:163–166, 1992.

P.H.A. Sneath and R.R. Sokal. *Numerical Taxonomy.* W.H. Freeman, San Francisco, 1973.

W.R. Softkey and C. Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *Journal of Neuroscience*, 13:334–350, 1993.

R.R. Sokal and F.J. Rohlf. *Biometry.* W.H.Freeman, New York, 1995.

M. Stengl, G. Ziegelberger, I. Boekhoff, and J. Krieger. Perireceptor events and transduction mechanisms in insect olfaction. In B.S. Hansson, editor, *Insect Olfaction*, pages 49–66. Springer-Verlag, 1999.

G. Stuart, N. Spruston, and M. Häusser. *Dendrites.* Oxford University Press, Oxford, 1999.

X.J. Sun, L.P. Tolbert, and J.G. Hildebrand. Synaptic organization of the uniglomerular projection neurons of the antennal lobe of the moth manduca sexta: A laser scanning confocal and electron microscope study. *Journal of Comparative Neurology*, 379:2–20, 1997.

Y. Tamaki and T. Yushima. Sex pheromone of the cotton leafworm, spodoptera littoralis. *Journal of Insect Physiology*, 20:1005–1014, 1974.

F.E. Theunissen and J.P. Miller. Temporal encoding in nervous systems - a rigorous definition. *Journal of Computational Neuroscience*, 2:149–162, 1995.

J. L. Todd, S. Anton, B. S. Hansson, and T. C. Baker. Functional organization of the macroglomerular complex related to behaviourally expressed olfactory redundancy in male cabbage looper moth. *Physiological Entomology*, 20:349–361, 1995.

N. Ulanovsky, L. Las, and I. Nelken. Processing of low-probability sounds by cortical neurons. *Nature Neuroscience*, 6:391–398, 2003.

P.G. Valeur, B.S. Hansson, and C. Löfstedt. Real-time measurement of pheromone release from individual female moths and synthetic dispensers in a wind tunnel by recording of single receptor-neurone responses. *Physiological Entomology*, 24:240–250, 1999.

P.G. Valeur, B.S. Hansson, K. Markebo, and C. Löfstedt. Relationship between sex pheromone elicited behaviour and response of single olfactory receptor neurons in a wind tunnel. *Physiological Entomology*, 25:223–232, 2000.

M.J. Van den Berg and G. Ziegelberger. On the function of the pheromone binding protein in the olfactory hairs of antheraea polyphemus. *Journal of Insect Physiology*, 37:79–85, 1991.

M.C.W. van Rossum and R.G. Smith. Noise removal at the rod synapse of mammalian retina. *Visual Neuroscience*, 15:809–821, 1998.

N.J. Vickers and T.C. Baker. Flight of heliothis virescens males in the field in response to sex pheromone. *Physiological Entomology*, 22:277–285, 1997.

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

J.D. Victor and K.P. Purpura. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology*, 76:1310–1326, 1996.

J.D. Victor and K.P. Purpura. Metric-space analysis of spike trains: theory, algorithms and application. *Network: Computational Systems*, 8:127–164, 1997.

J. von Neumann. Various techniques used in connection with random digits. In *Collected Works, Volume 5.*, pages 768–770. Pergamon Press, New York, 1963.

M. Wachowiak, W. Denk, and R.W. Friedrich. Functional organization of sensory input to the olfactory bulb glomerulus analyzed by two-photon calcium imaging. *Proceedings of the National Academy of Sciences*, 101:9097–9102, 2004.

J. Walraven, C. Enroth-Cugell, D.C. Hood, M.L. Dia, and J.L. Schnapf. The control of visual sensitivity. In L. Spillmann and S.J. Werner, editors, *Visual perception: the neurophysiological foundations.*, pages 52–101. Academic Press, San Diego, 1990.

D.S. Wei, Y.A. Mei, A. Bagal, J.P. Kao, S.M. Thompson, and C.M. Tang. Compartmentalized and binary behavior of terminal dendrites in hippocampal pyramidal neurons. *Science*, 293:2272–2275, 2001.

M.A. Willis and T.C. Baker. Paper title. *Effects of intermittent and continuous pheromone stimulation on the flight behaviour of the oriental fruit moth, Grapholita molesta.*, 9:341–358, 1984.

E.O. Wilson. Pheromones. *Scientific American*, 208:100–114, 1963.

P. Witzgall, T. Lindblom, M. Bengtsson, and M. Tóth. The Pherolist. www-pherolist.slu.se, 2004. (page www-pherolist.slu.se/pherolist.php?command=phero list&file=compound _index&id=87&ordered=1000 accessed on 2nd May 2006).

J.M. Zanker and J.P. Harris. On temporal hyperacuity in the human visual system. *Vision Research*, 42:2499–2508, 2002.

F. Zufall and T. Leinders-Zufall. The cellular and molecular basis of odor adaptation. *Chemical Senses*, 25:473–481, 2000.

Appendix A

Principal Components Analysis

A.1 Overview

The technique of principal components analysis (PCA) has been used in this thesis in two different areas. Firstly, in Chapter 2, PCA was used as part of the spike sorting procedure to find which aspects of a spike shape can be used to separate out the noisy signal from spikes generated by several neurons (Section 2.3.2.1). Secondly, in Chapter 4, PCA was used as part of the procedure to find whether the aspects of a binned spike train showing the most variation are correlated with the stimulus or not (Section 4.4.1). The purpose of this appendix is to clarify the method used to find the principal components.

A.2 PCA Description

In data sets with many variables, it is often possible to find that a single driving factor of the system may be measured by more than one variable. This redundancy of information allows for simplification of the problem by replacing a group of variables by a single new variable. This simplification can be achieved using PCA, where each PC is a linear combination of the original variables, and since each PC is orthogonal to all other PCs, there is no longer any redundancy.

The PCA procedure is such that the projection of each observation onto the first PC (a single axis in space), results in a new variable, the variance of which is the maximum

among all possible choices of the first PC. The projection of each observation onto the second PC (another axis in space, perpendicular to the first), results in another new variable, the variance of which is the maximum among all possible choices of the second PC. While the full set of PCs is the same size as the original data set, a high proportion of the total variance of the original data is now explained by the first few new variables.

A.3 PCA Calculation

Throughout this thesis, PC's have been calculated using MATLAB's `princomp` function, which involves performing singular value decomposition on the normalised covariance matrix of the input matrix. For the input matrix, $\mathbf{X}$, (where each row is an observation, and each column is a variable), the covariance matrix, $\mathbf{C}$, can be calculated as $\mathbf{C} = \mathbf{X}\mathbf{X}^{\mathrm{T}}/(n-1)$, where n is the number of observations. The singular value decomposition is then performed on $\mathbf{C}$, where the singular value, σ, and corresponding singular vectors, $\mathbf{u}$ and $\mathbf{v}$, satisfy $\mathbf{C}\mathbf{v} = \sigma\mathbf{u}$ and $\mathbf{C}^{\mathrm{T}}\mathbf{u} = \sigma\mathbf{v}$. With the singular values on the diagonal of a diagonal matrix, $\mathbf{\Sigma}$, and the corresponding singular vectors forming the columns of two orthogonal matrices $\mathbf{U}$ and $\mathbf{V}$, the singular value decomposition is $\mathbf{C} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathrm{T}}$.

Each singular vector is used to calculate the linear combinations of the original variables to give the PC scores, and each singular value indicates the variance accounted for by each PC. The singular values and vectors are then used to transform the original data set by projecting the original points onto the new axes. When $\mathbf{C}$ is square, symmetric and positive definite, the singular value decomposition is equivalent to the eigenvalue decomposition.

A.4 Further Reading

For further information, see Cooley and Lohnes, 1971; Everitt and Dunn, 1999; Jackson, 1988; Krzanowski, 1988; Seber, 1984.

Appendix B

Inhibited ALN Analysis

B.1 Overview

Chapter 3 includes a short analysis involving an ALN showing inhibition to the presented pheromone stimulus. The method used to calculate the dose response curve involved the instantaneous spike frequency (Section 3.4.5.2). Later consideration, described below, shows that the net spikes s^{-1} could also have been used to create the dose response with similar outcome.

B.2 Modified Net Spikes Method

In order to calculate the net spikes s^{-1} due to the stimulus, where the response is a reduction in firing rate, a similar procedure to that described in Section 3.3.2 can be used. The only difference being that the start and end of the time period, t, were taken to be the first of three consecutive time bins where the firing rate rose above the threshold value (set to the mean background firing rate minus three standard deviations). The dose response curve created using this method suggests that the ALN responds to the stimulus with a suppression of the background firing rate at all stimulus loads tested (Figure B.1).

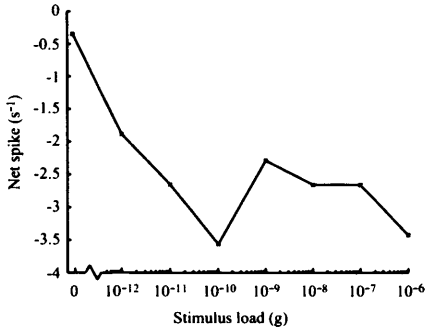

Figure B.1: Dose Response Curve of Inhibited Antennal Lobe Neuron using the Net Spikes Method. With a slightly modified version of the net spikes s^{-1} calculation, the dose response curve of the inhibited ALN from Figure 3.7 shows fewer net spikes s^{-1} in response to all stimulus loads tested compared to that of the blank stimulus.

Appendix C

Effects of Variation of Parameters on Stationarity Test

C.1 Overview

When designing the test for stationarity described in Section 4.3.6, two main parameters could be varied, the length of the window used for the moving average filter, and the number of standard deviations used to define the upper and lower bounds. While making little difference to the final outcome, variations in these parameters are shown here for completeness.

C.2 Variation of Moving Average Filter Window Size

The stationarity test performed in Section 4.4.4 used a window size equivalent to 0.5 seconds. Window sizes equivalent to 0.25 seconds and 1 second were also tested. All three window sizes are presented here, for both 5 ms and 10 ms time bin sizes, showing numbers of recordings found to be stationary and non-stationary (Tables C.2, C.3 and C.4). For ease of comparison, a summary table shows the percentage of stationary recordings at each stimulus load using all three window durations (Table C.1).

The timing of the stimulus onset for the real neuronal data itself influenced the final decision on window size. Most ORNs had a stimulus onset time of at least 1 second into the recording (some had 2 seconds before the stimulus onset). Using a window duration

a) Bin size 5 ms	% Stationary at stimulus loads:				
Window Duration	0	10^{-9} g	10^{-8} g	10^{-7} g	Total
0.25 s	41%	40%	16%	27%	31%
0.5 s	59%	40%	25%	32%	39%
1 s	50%	40%	41%	30%	40%

b) Bin size 10 ms	% Stationary at stimulus loads:				
Window Duration	0	10^{-9} g	10^{-8} g	10^{-7} g	Total
0.25 s	44%	46%	16%	24%	33%
0.5 s	57%	42%	28%	31%	40%
1 s	50%	43%	42%	25%	40%

Table C.1: Summary of Effect of Moving Average Window Duration. Using boundaries of $\mu_{MA} \pm (3 \times \sigma_{MA})$, the use of moving average filters with different window durations causes some modifications of numbers of stationary spike trains across the individual stimulus loads. In particular, at the two higher stimulus loads, the longer the window duration, the more recordings are deemed stationary. At the two lower stimulus loads, and over the total number of recordings, the window duration shows little effect.

a) Bin size 5 ms	Stimulus loads					
Window 0.25 s	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	29	26	10	19	84	31%
Non-stationary	41	39	54	52	186	69%
Total	70	65	64	71	270	
%Stationary	41%	40%	16%	27%	31%	

b) Bin size 10 ms	Stimulus loads					
Window 0.25 s	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	31	30	10	17	88	33%
Non-stationary	39	35	54	54	182	67%
Total	70	65	64	71	27	
%Stationary	44%	46%	16%	24%	33%	

Table C.2: Moving Average Window Size = 0.25 seconds

equivalent to either 0.25 seconds or 0.5 seconds allowed the initial start up transient of the filter to die away before the mean value of the filtered background firing rate was calculated. For a window duration of 1 second, the mean of the background firing rate was calculated from an end section of recording after the response is assumed to have died down. In a small number of cases, the ORNs were still showing a response to the stimulus presentation at the end of the recording. The window duration of 0.5 seconds was chosen for the main analysis

a) Bin size 5 ms	Stimulus loads					
Window 0.5 s	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	41	26	16	23	106	39%
Non-stationary	29	39	48	48	164	61%
Total	70	65	64	71	270	
%Stationary	59%	40%	25%	32%	39%	

b) Bin size 10 ms	Stimulus loads					
Window 0.5 s	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	40	27	18	22	107	40%
Non-stationary	30	38	46	49	163	60%
Total	70	65	64	71	27	
%Stationary	57%	42%	28%	31%	40%	

Table C.3: Moving Average Window Size = 0.5 seconds

a) Bin size 5 ms	Stimulus loads					
Window 1 s	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	35	26	26	21	108	40%
Non-stationary	35	39	38	50	162	60%
Total	70	65	64	71	270	
%Stationary	50%	40%	41%	30%	40%	

b) Bin size 10 ms	Stimulus loads					
Window 1 s	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	35	28	27	18	108	40%
Non-stationary	35	37	37	53	162	60%
Total	70	65	64	71	27	
%Stationary	50%	43%	42%	25%	40%	

Table C.4: Moving Average Window Size = 1 second

since it also gave a smoother representation of the neuronal firing rate for the duration of the recording.

C.3 Variation of Number of Standard Deviations

The stationarity test performed in Section 4.4.4 used boundaries based on $\mu_{MA} \pm (3 \times \sigma_{MA})$. Boundaries based on 2 and 4 standard deviations were also tested. All three values are presented here, for both 5 ms and 10 ms time bin sizes, showing numbers of recordings found

to be stationary and non-stationary (Tables C.3, C.6 and C.7). For ease of comparison, a summary table shows the percentage of stationary recordings at each stimulus load using all three numbers of standard deviations (Table C.5).

a) Bin size 5 ms	% Stationary at stimulus loads:				
# of σ_{MA}	0	10^{-9} g	10^{-8} g	10^{-7} g	Total
2	31%	20%	22%	13%	21%
3	59%	40%	25%	32%	39%
4	64%	58%	45%	45%	53%

b) Bin size 10 ms	% Stationary at stimulus loads:				
# of σ_{MA}	0	10^{-9} g	10^{-8} g	10^{-7} g	Total
2	30%	22%	20%	14%	21%
3	57%	42%	28%	31%	40%
4	63%	58%	41%	45%	52%

Table C.5: Summary of Effect of Number of Standard Deviations. Using a window duration of 0.5 s, the number of standard deviations used to classify recordings as stationary or not cause differences across the individual stimulus loads, and in the total number. However, even with the largest number of standard deviations tested, just over 50% of the recordings were deemed stationary.

a) Bin size 5 ms	Stimulus loads					
$2 \times \sigma_{MA}$	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	22	13	14	9	58	21%
Non-stationary	48	52	50	62	212	79%
Total	70	65	64	71	270	
%Stationary	31%	20%	22%	13%	21%	

b) Bin size 10 ms	Stimulus loads					
$2 \times \sigma_{MA}$	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	21	14	13	10	58	21%
Non-stationary	49	51	51	61	212	79%
Total	70	65	64	71	27	
%Stationary	30%	22%	20%	14%	21%	

Table C.6: Boundary using $2 \times \sigma_{MA}$

For the UEA method to be meaningful, stationarity of recordings is essential (Section 4.3.4), so any of the thresholds tested here could have been used in the final analysis. The main point of the this test is to imply that not all recordings are stationary, and therefore

a) Bin size 5 ms	Stimulus loads					
$4 \times \sigma_{MA}$	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	45	38	29	32	144	53%
Non-stationary	25	27	35	39	126	47%
Total	70	65	64	71	270	
%Stationary	64%	58%	45%	45%	53%	

b) Bin size 10 ms	Stimulus loads					
$4 \times \sigma_{MA}$	0	10^{-9} g	10^{-8} g	10^{-7} g	Total	%
Stationary	44	38	26	32	140	52%
Non-stationary	26	27	38	39	130	48%
Total	70	65	64	71	27	
%Stationary	63%	58%	41%	45%	52%	

Table C.7: Boundary using $4 \times \sigma_{MA}$

the stationarity assumption is violated.

Appendix D

Investigation into Spike Patterns

D.1 Overview

When investigating the possible structure of temporal encoding in individual ORNs at the periphery, the UEA procedure (Grün, 1996; Grün et al., 2001a) was first employed (Section 4.4.3). Unfortunately, the ORN data set available shows a non-stationary firing rate across the duration of many of the recordings (Section 4.4.4), violating a vital assumption of the UEA method (Section 4.3.4). A further investigation of the synchrony of spikes across the population of available ORNs using the UEMWA procedure (Grün et al., 2001b) encountered problems with maintaining a high enough number of spikes for statistical viability whilst preserving the stationarity of spike train sections (Section 4.4.5). Another approach to the investigation of temporal patterns is perhaps to investigate the similarities between spike trains rather than searching specifically for patterns. Fellous et al. (2004) suggest such a method, in which both the stationarity and the firing rate of spike trains are not an issue. An initial investigation into the use of this method is presented here.

D.2 Discovering Spike Patterns

This section gives a brief explanation of the method used to discover spike patterns, as detailed by Fellous et al. (2004). The basic idea is that by using repeat presentations of a

stimulus, patterns of spikes can be formed by neurons in response to these presentations, but that although reliable, the actual structure of the spike pattern can vary according to the pre-stimulus history of the neuron. The different patterns of spikes can be better visualised by reordering raster plots, from multiple repeat recordings that are time-locked to the onset of a stimulus presentation, according to the similarity between the individual spike trains.

For this particular analysis, the times of the spikes are first found from the voltage traces (see Chapter 2), resulting in a sequence of δ functions, which are then smoothed by convolving with a Gaussian kernel of fixed width σ. From these convolved spike trains, a symmetrical similarity matrix, $\mathbf{S}$ (see Figure D.1 b) for an example) is formed by calculating the normalised dot product between all pairs of trials $< i, j >$:

$$s_{ij} = \frac{\mathbf{g}_i \cdot \mathbf{g}_j}{\|\mathbf{g}_i\| \cdot \|\mathbf{g}_j\|} \tag{D.1}$$

where $\mathbf{g}_i$ represent the convolution of spike train i with the Gaussian kernel. Values of s_{ij} lie between 0 and 1 resulting in a measure of the similarity or correlation between the input vectors $\mathbf{g}_i$ and $\mathbf{g}_j$.

In general, values within the similarity matrix were confined to a small region of the overall range (see Figure D.1 c)). To increase the range over which the similarity values lay, all values within the similarity matrix were scaled using a sigmoid function:

$$b_{ij} = \frac{1}{1 + \exp^{-\frac{\left(s_{ij} - \bar{s_u}\right)}{\tau}}} \tag{D.2}$$

The sigmoid was centred at the mean of the similarity values, with the slope, τ determined empirically:

i. The similarity values were collected into a histogram containing 50 bins between the minimum and maximum possible values of 0 and 1 respectively.

ii. The value of τ was incremented by 0.005 (from 0.01 to 0.3) and the similarity values s_{ij} were scaled using Equation D.2 to give new similarity values b_{ij}.

iii. The new values b_{ij} were then collected into a new histogram (using the same 50 bins) and the standard deviation among all the bins was calculated.

iv. The value of τ yielding the minimum standard deviation (i.e. the flattest histogram) was chosen as the final value with the search stopping if the smallest bin contained no values (i.e. totally dissimilar values no longer existed).

An N-dimensional vector of the scaled similarities between trial j and all other trials $i = 1, \ldots, N$ now represents each of the N spike trains (columns of the scaled similarity matrix). Using this recoding, trials with similar spike patterns have vectors that are close in the Euclidean sense, and can be grouped using a clustering algorithm.

Fellous et al. (2004) suggest using a fuzzy K-means algorithm for the clustering. This is a variant on the standard K-means algorithm where the strict membership condition of each point being assigned to only one cluster, is replaced by a probability-like membership of belonging to each of the K clusters (see Bezdek, 1981; Dumitrescu et al., 2000; Hoppner et al., 1999 for further explanation). The scaled similarity matrix, and hence the individual spike trains, were then reordered according to the output of the clustering algorithm (see Figure D.1).

D.3 Initial Test Using Surrogate Data

To test the implementation of the algorithm, a surrogate data set was used to ensure the output from the clustering algorithm performed as expected. Fellous et al. (2004) created the surrogate data sets used here, and these are available from http://www.cnl.salk.edu/~fellous/data/JN2004data/data.html. The data sets were created to contain a known number of clusters. Each cluster contained between 4 and 6 randomly chosen events around which spikes occurred. Each cluster was parameterised by three variables: the percentage of missing spikes across all trials; the number of extra non-event related spikes per trial;

and the average standard deviation of the spike time jitter around each event. The algorithm implemented for this appendix requires only the spike times as input, and these were provided with trial order randomised so as to remove the cluster structure (see Figure D.1 a)). The initial similarity matrix was then generated using Equation D.1 (Figure D.1 b)), with the histogram created from the similarity values shown in Figure D.1 c). This histogram was scaled using the sigmoidal function (Equation D.2) to produce the "flattest" histogram (Figure D.1 d)). The scaling was then applied to the similarity matrix (Figure D.1 e)) producing the input to the clustering algorithm. Input trials were then reordered according to the output of the clustering algorithm resulting in similar spike trains being grouped together. Timing events around which the initial spike times were grouped are now apparent (Figure D.1 f)).

D.4 Initial Test Using ORN Data

An initial test was performed on repeat recordings at a single stimulus load from a single ORN. For this example case, there are 56 available recordings on which to perform the analysis. For this analysis, two different values of σ, the length of the Gaussian kernel, were used. Firstly, with $\sigma = 20$ ms, the result of the reordering from the clustering output shows the recordings arranged such that the noticeable difference between clusters is the time at which the neuron starts to respond to the stimulus (Figure D.2). This difference in response onset time appears to be a major feature of the recordings (combined with the large increase in firing rate), and could be due to changes with the stimulus delivery system, or with the electrode location. Secondly, with $\sigma = 5$ ms, the result of the reordering shows similar structure in the similarity matrix and the raster plot (Figure D.3).

The value of σ plays an important role in deducing the similarity of the recordings, and thus the clustering algorithm. When σ is large, the emphasis is on the overall firing rate of the recordings, whilst when σ is small, the emphasis is on the finer temporal structure

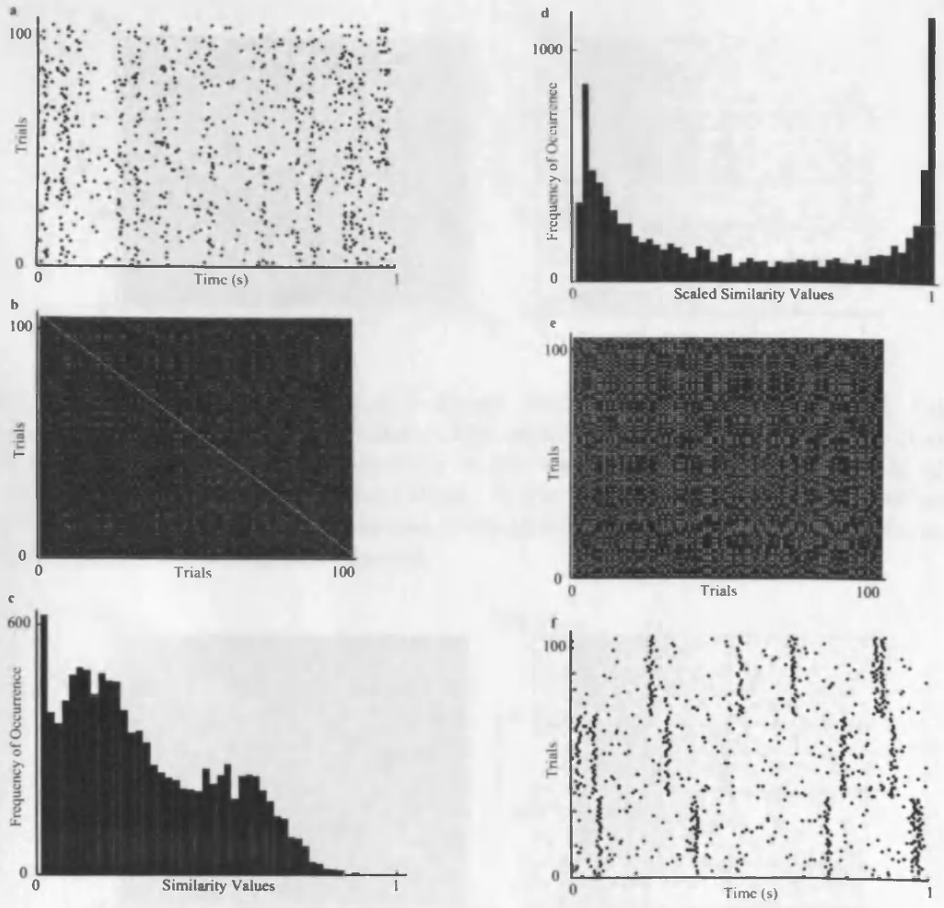

Figure D.1: Initial Test of Spike Patterns Implementation. a) Input Spike Trains. A surrogate data set from http://www.cnl.salk.edu/~fellous/data/JN2004data/data.html, was used to test the implementation of the clustering algorithm developed by Fellous et al. (2004). The ordering of 105 repeated trials, containing several spike events, was randomised and when the raw spike times are displayed as a raster plot there is little obvious structure. b) Similarity Matrix. Input spike trains were convolved with a Gaussian kernel. The similarity matrix was then generated using Equation D.1, with lighter shades representing high similarity values. c) Similarity Histogram. Similarity values were then split into 50 bins between 0 (dissimilar spike trains) to 1 (similar spike trains), resulting in a distribution usually centred below 0.5. d) Scaled Similarity Histogram. A sigmoid function was employed to augment the range of similarity values (Equation D.2). The slope of the sigmoid was chosen such that the scaled histogram had the lowest standard deviation among all bins (i.e. the "flattest" histogram) with the constraint that the smallest bin contained values (i.e. dissimilar values existed). e) Scaled Similarity Matrix. The scaling was then applied to the similarity matrix resulting in the scaled similarity matrix, the input to the clustering algorithm. f) Reordered Trials. Input trials were then reordered according to the output of the clustering algorithm. The random events around which some of the spikes were distributed are now more apparent visually.

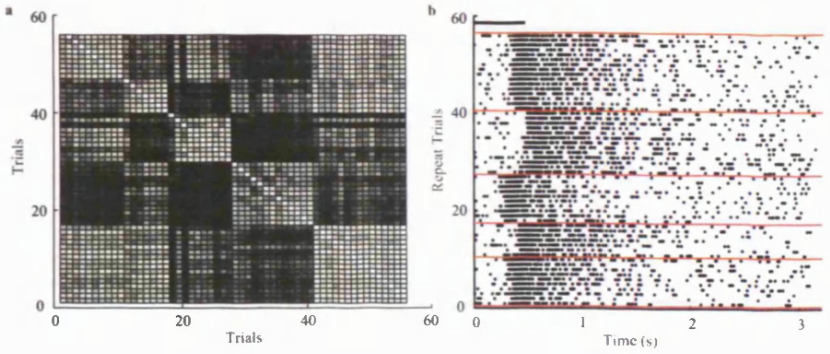

Figure D.2: Test Using ORN Data, $\sigma = 20$ ms. a) Reordered Similarity Matrix. Lighter shades represent higher similarity values. The reordering suggested by the output of the clustering algorithm introduces structures in the similarity matrix corresponding to the five clusters used. b) Reordered Recordings. Raster plot shows clusters arranged mainly according to the onset time of the response to the stimulus (horizontal black bar). Horizontal red bars separate recordings into clusters.

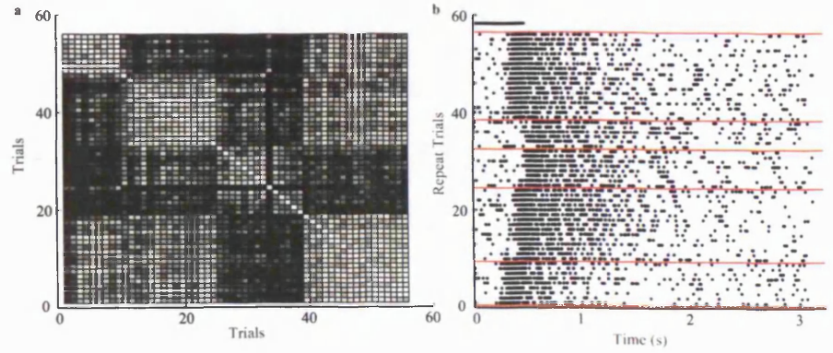

Figure D.3: Test Using ORN Data, $\sigma = 5$ ms. a) Reordered Similarity Matrix. Lighter shades represent higher similarity values. The reordering suggested by the output of the clustering algorithm again introduces structures in the similarity matrix corresponding to the five clusters used. b) Reordered Recordings. Raster plot shows clusters again arranged according to the onset time of the response to the stimulus (horizontal black bar). Horizontal red bars separate recordings into clusters.

(Figure D.4). Fellous et al. (2004) suggest that when σ was approximately equal to the amount of jitter in the data the result was optimal.

Further analyses using this method have yet to be performed. One main constraint is the amount of data required. The present data set contains only a maximum of five recordings at each of the stimulus loads for most ORNs, but this method of finding patterns in spikes requires many more.

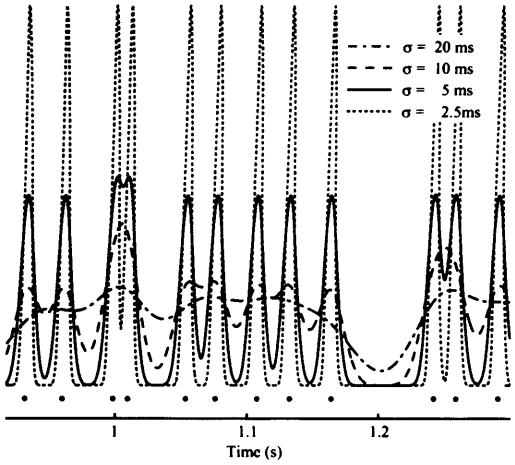

Figure D.4: Effect of Changing σ. A section of a recording is shown with a range of σ values used in the Gaussian convolution operation. Larger values of σ act as a low pass filter, smoothing the effect of each spike. Smaller values of σ emphasise individual spikes. Black dots at the bottom represent the actual spike times.