

Olfactory Object Recognition and Generalization in the Insect Brain

Thesis by
KAI SHEN

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



California Institute of Technology
Pasadena, California
2010
(Defended March 18, 2010)

© 2010

Kai Shen

All Rights Reserved

Acknowledgment

First of all, I would like to thank my graduate advisor, Gilles Laurent for leading me into a fascinating area of research; for both his experimental rigor and clarity of thought; for his appreciation of beautiful data; for his ability to ask the right questions and weave together a great story. I am also very grateful for his understanding during a difficult period in my graduate career, and for allowing me to take a sabbatical to explore other interests, and finally for providing the means for me to travel to Germany to write my thesis. I would also like to thank the other members of my thesis committee: Erik Winfree, for his insights, and valuable suggestions on data analysis; Thanos Siapas, for his skepticism and methodical approach to science that is unsurpassed, and for providing tetrodes for KC recordings that became the foundation to this thesis; and Christof Koch and Michael Dickinson for their valuable support over the years.

I would like to thank Sina Tootoonia, who I collaborated with on the PN-KC model described in Chapter 2; Sina built the model in C. I have truly enjoyed our conversations about science and otherwise, and I am grateful for his generosity, incredible work ethic and dedication to rigor. I would also like to thank Anusha Narayan and Maria Papadopoulou for valuable contributions to the PN-KC model. From the time I arrived in the Laurent lab until now, I have greatly benefited from the generosity of others, I would like to thank especially: Stjin Cassenaer for his key insights and

advice, and extraordinary experimental prowess that is inspiring; Glenn Turner, for always being a willing listener, and for his good humor and for his love of science; Cindy Chiu for her encouragement and thoughtfulness; Ben Rubin for all the basketball and life conversations; Vivek Jayaraman for helping me get started in the early days and helpful criticisms; Appreciation also go out to Ofer Mazor, Bede Broome and Mark Stopfer for their advice and expertise in KC recordings; Ueli Rutishauser for discussions about RLSC and science in general; Thomas Nowotney and Ramon Huerta for interesting ideas and different ways of thinking about locust olfaction; and Viola Priesemann for her help in the writing process.

On a personal level, I would like thank Laurent Demanet, Zhang Cheng Zhong, Adrian Noetzli, Teresa Wong and Caroline Chiang for their friendship, support and encouragement over the years. To the folks at Prufrock house, past and present, especially Lily Liu, Zuli Kurji, Matthew Kelley and Jasper Simon for all the fun memories. To all the kids that came out to volunteer at Union Station and to the Caltech Y, – there is really no organization quite like it. To Alice Lin, for staying up late to help me format my thesis and making last minute figure legends, and for bringing light and color into my life. To my brother Hong, for always believing in me. And to my Mum and Dad, for making sacrifices in their lives so that mine could be better.

Abstract

Sensory object recognition is the most fundamental of operations performed by the brain. A key computational difficulty of object recognition is that it requires both selectivity to particular objects (e.g., exact odor mixture identification) and generalization across objects (identifying particular features or components common to different odors). Although previous results (*1*) suggest that odor identity and intensity are represented in the activity of both PNs and KCs, it is not clear how these representations generalize across complex odor mixtures. In particular, it is not clear what types of information are available in KC population (or if its even possible to decode across KC populations?) and how is this information represented? Using the locust olfactory system as a model system, we found that Kenyon cells (KCs), the principal neurons of the mushroom body, an area required for associative learning can identify the presence of components in mixtures and thus enable odor segmentation. As a population, small groups of KCs can both identify and categorize odors with high accuracy. We identified and tested simple circuit requirements for this computation, and propose that odor representations in mushroom bodies are optimized for odor memorization, identification and generalization. These rules may be relevant for pattern classifying circuits in general.

Contents

	Page
ACKNOWLEDGMENTS	iii
ABSTRACT	v
TABLE OF CONTENTS	vi
1. INTRODUCTION	1
1.1 Object Recognition	1
1.2 Olfactory Computations	5
<i>1.2.1 Nature of odors</i>	
<i>1.2.2 Perception of odor mixtures</i>	
<i>1.2.3 Odor discrimination</i>	
<i>1.2.4 Odor segmentation</i>	
<i>1.2.5 Odor generalization</i>	
1.3 Olfactory Microcircuits	12
<i>1.3.1 Olfactory receptor neurons</i>	
<i>1.3.2 Antennal lobe</i>	
<i>1.3.3 Antennal lobe-to-Mushroom body circuit</i>	
1.4 Outline and Specific Aims	18
2. MIXTURE CODING AND ODOR-SEGMENTING KENYON CELLS	20
2.1 Results	22
<i>2.1.1 Representation of binary mixtures by PNs</i>	

2.1.2	<i>Representation of mixtures of increasing complexity by PNs</i>	
2.1.3	<i>Kenyon cell responses to mixtures</i>	
2.1.4	<i>PN and KC population statistics</i>	
2.1.5	<i>Generating model KC classifiers using recorded PN data</i>	
2.2	Discussion	42
2.2.1	<i>Functional consequences</i>	
2.2.2	<i>KCs as independent feature detectors?</i>	
2.3	Methods	75
2.3.1	<i>Preparation and stimuli</i>	
2.3.2	<i>Binary-mixture experiments</i>	
2.3.3	<i>Multi-component-mixture experiments</i>	
2.3.4	<i>Electrophysiology</i>	
2.3.5	<i>Recording constraints and potential sampling biases</i>	
2.3.6	<i>Extracellular data analysis</i>	
2.3.7	<i>Analysis</i>	
2.4	Acknowledgements	96
3.	ODOR IDENTIFICATION AND GENERALIZATION IN THE MUSHROOM BODY	97
3.1	Results	98
3.1.1	<i>Decoding PN trajectories over time</i>	
3.1.2	<i>Decoding odor identity and category from population KC vectors</i>	

3.2 Discussion	105
<i>3.2.1 Odor segmenting KCs leads to stimulus generalization</i>	
<i>3.2.2 Combinatorial codes for mixtures in the MB</i>	
3.3 Methods	126
<i>3.3.1 Mean KC spike latency</i>	
<i>3.3.2 Population decoding</i>	
4. SUMMARY AND CONCLUSIONS	132
4.1 An elegant solution for odor recognition and generalization	132
4.2 Open questions	134
BIBLIOGRAPHY	137

CHAPTER 1

Introduction

Sensory object recognition is the most fundamental of operations performed by the brain. Computational vision, e.g., has been directed at understanding how visual objects and their associated properties are reconstructed and recognized from local information that falls on the retina (1). In the olfactory system, object recognition involves identifying characteristic combinations of molecules in complex blends of odorants. Three fundamental computations olfaction must solve are: recognition, concentration-invariance, and mixture segmentation.

1.1 Object Recognition

The apparent ease with which we recognize objects belies the magnitude of this feat: we effortlessly recognize objects from among millions of possibilities and we do so within a fraction of a second, in spite of tremendous variation in the appearance of each one (2). Understanding the computational processes that underlie this ability is one of the fundamental goals of neuroscience. Object recognition can be defined as

the ability to accurately discriminate each object from every other object ('identification') or set of similar objects ('categorization') from all other possible objects, and to do this over a range of identity preserving transformations (e.g., variation in object position, scale, pose and illumination, and the presence of visual clutter) (2-4). We see an object many times but never see the exact same image on our retina twice (5). Although several efforts have examined this so called "invariance problem" (6-8), a robust, recognition machine still evades us and we lack a truly satisfying understanding of how the problem is solved by the brain.

In the visual cortex, evidence suggest that neuronal processes underlying object recognition are localized in the ventral stream (V1→V2→V4→posterior IT→anterior IT) (9). The highest stages of this stream are thought to convey neuronal signals that explicitly support object recognition - i.e., object identity can be directly extracted from these populations despite identity-preserving image changes (4, 10, 11). How are such useful neuronal representations constructed in the brain? How is it that neurons can be sensitive to subtle changes in object identity, and yet are relatively insensitive to large identity-preserving image changes?

Olfaction is very different to vision, because it's a synthetic sense (12). However, many of the recognition problems are still the same.

Hopfield (13) postulated several computations that olfactory systems need to solve that are more sophisticated than simple discrimination. These include: (i) concentration-invariant odor recognition, (ii) background-invariant odor recognition, and (iii) odor-mixture segmentation. Such computations could underlie complex behaviors, e.g., by tracing a series of sequential odor way-points, salmon can navigate over thousands of miles from their spawning grounds to the open ocean and back (14), and rodents have also been shown to track odors (15). These are examples of computations that have been proposed but remain poorly characterized at the behavioral level and unexamined at the neural level.

To recognize an object, an animal must use some internal representation of the external stimulus (e.g., visual or olfactory scene) to make a decision: is object *A* present or not? Computationally, the brain must apply a decision function to divide an underlying neuronal representational space into regions where object *A* is present and regions where it is not (16, 17). Because brains are essentially networks of neurons, the subject must also have read-out neurons that can successfully report if object *A* was present or not. The central computational issues of object recognition are (2):

(i) What is the format of the representation used to support the decision (the substrate on which the decision functions directly operate)?

i.e., how is the external world represented and what information is available in the neural code?

(ii) What kinds of decision functions (i.e., read-out tools) are applied to that representation? i.e., how is the neural code read out by decision neurons?

Object recognition can thus be viewed as the problem of finding operations that progressively transform the external world (e.g., visual, olfactory, auditory stimuli) into a new form of representation, followed by the application of relatively simple decision functions (e.g., linear classifiers) that ultimately lead to perception and behavior (2). Along the way, there are other issues, such as how many neurons are required in computing the decision function, where are they in the brain, is their operation fixed or dynamically related to the task (i.e., nonlinearities involved, learning) and how they code their choices (as a function of spikes)?

Basic machine learning textbooks tell us that feature selection is more important than the complexity of the classifier used (SVM, NN or one of many nonlinear classifiers). It has been shown that a variety of recognition tasks can be solved in inferotemporal (IT) cortex using simple linear classifiers (10, 18). The feature vector that results from the preprocessing stage of external stimuli is oftentimes not suitable to be

applied by a decision function due to its high dimensionality and redundancy. The “curse of dimensionality” tell us that the number of training examples must grow exponentially with the number of features in order to learn an accurate model (19). Because only a limited number of examples are typically available, there is an optimal number of feature dimensions beyond which the performance of the recognition model starts to degrade. Thus, dimensionality reduction is crucial in order to build a classifier with high accuracy. We can think of progressive transformations from one form of representation to another as extracting features optimal for decision making, akin to dimensionality reduction.

1.2 Olfactory Computations

1.2.1 Nature of odors

The sense of smell is well known for its ability to detect odorants at levels far below that of the most sensitive instruments, to discriminate between thousands of single odorants, and to engrave into memory recollections of smells that stretch back in time. Our olfactory abilities to detect, discriminate, and imprint odors is unmatched by any electronic nose. Naturally occurring odors rarely emanate from a single chemical, it consists of hundreds to thousands of components, of varying ratios. The

sense of smell must have developed, over evolution, to process complex odor mixtures.

The relationship between odor structure, neural representations and perception is not deterministic, and depends on experience and learning (20). Odors are also difficult to quantify because its chemical structure cannot be described by a simple set of parameters – odor space is extremely high dimensional and vary in an indeterminate number of parameters such as carbon–chain length, molecular weight, and polarity. Recently a multidimensional, physicochemical odor space was devised to describe odor structure: 1664 molecular descriptors (thus 1664–d space) for 1500 odors (21). Each odor was represented as a 1664–d vector in this space. This odor space maybe useful in predicting odor perception. Principle component analysis revealed a correlation between odor structure in this space and its perceived pleasantness among humans (22).

1.2.2 Perception of odor mixtures

Over the centuries, odor mixtures in the form of fragrances and incense have been used to suppress or mask a variety of unpleasant odors, including body odors or those emanating from waste materials, as well as to attract mates of the opposite sex. Through trial and experience, the

most liked and effective fragrance mixtures became known, - e.g., jasmine, rose and orange oils are still in use. Our knowledge of brain mechanisms that computes and recognizes odor mixtures is however very limited. The study of the perception of odor mixtures by humans starts with binary mixtures, e.g., (23), and there are several common outcomes: both odors may be identified with no change in their perceived intensities compared to their intensities before mixing; both may be perceived but with one or both having reduced intensity; or one may be suppressed to such an extent that it cannot be perceived (24-27). Odor reduction or suppression is dependent on the particular odor combination and the concentration of each odor. When binary mixtures are analyzed by human subjects, discrimination of the individual components is influenced not only by the type of odor and the perceived intensity of each component, but also by their familiarity and pleasantness (28).

In humans, correctly judging whether a binary mixture is different (discrimination) from a stimulus consisting of only one of the components is generally much simpler than having to identify one or both of the components (identification). The ability to identify components in mixtures becomes increasingly more difficult as the number of components increase, with the limit at ~4 (25-27). Furthermore, regardless of whether a choice was correct or incorrect, 4

was also the most common maximum number of odors chosen in mixtures containing up to 8 odors. When this experiment was repeated with mixtures instead of single components, the result was again 4. Taken together, this indicates that humans encode complex odor mixtures as single entities.

1.2.3 Odor discrimination

With training, rodents can learn to discriminate between virtually any pair of pure odors, including highly related stereoisomers (29–31). Rats can also perform difficult odor discrimination tasks after extensive lesions of the OB (29), suggesting that simple discrimination (e.g., go/no-go, where 1 bit of information to be extracted) is fundamental to olfactory processing and not necessarily a computationally difficult task (or at least not one that requires the entirety of the OB).

1.2.4 Odor segmentation

The 2007 Pixar film *Ratatouille* is about how a rat Remy, blessed with unusually sharp olfactory senses, becomes an extraordinary chef. The food odors in Remy's kitchen do not exist alone, there are experienced in the larger context of a variety of different odors and thoroughly mixed. For example, Remy must be able to identify by smell, different spices and

ingredients that goes into a soup. Odor segmentation is thus the ability to identify unique odor objects from a mixture of odors, – an analogy to the “cocktail party problem” in audition – where a single voice is recognized from a background of many different noises (32). Psychophysical experiments in humans show that the limit of identifying components within mixtures is ~ 4 (24), but if asked to identify only one highly familiar odor within mixtures, performance becomes chance only after 16 odors in the mixture (33). Highly olfactory animals such as rats likely can do even better (34).

Because odors are thought to be mixed thoroughly at the receptor level, the computational problem of segmenting particular components is not necessarily a simple one. By exploiting temporal correlations between groups of receptors to temporal fluctuations (in concentrations) of different odor streams, odor mixtures have been decomposed into its components in a computational model (35). However, the problem becomes more difficult if the temporal fluctuations are not present, as is the case when the different components are very well mixed (e.g., aroma from wine or well cooked meal). Brody and Hopfield (34) implemented a neural model (termed many-are-equal) using spike-timing computations allowing for concentration-invariant recognition and odor segmentation within complex olfactory scenes. Their model relies on differential phase of firing of principal cells to different odors. However, this is unlikely to

apply to the locust antennal lobe (AL, model system of this thesis), because in our model system, the phase of AL principal neuron firing contains no information about odor identity or intensity (7).

1.2.5 Odor generalization

Learning strongly influences the ability of animals to perceive and identify odors. Animals must take information about a conditioned odor experience and generalize this information to future odor experiences because no two conditioning stimuli are experienced in exactly the same way (36). Generalization of a conditioned response, (e.g., a learned association of an odor with reward), occurs when animals perceive similarities among stimuli from one experience to the next. Put it more formally, it is the ability to decide that two odors, though readily distinguishable, are similar enough to afford the same outcome (36). One beautiful example in nature is the foraging behavior by honeybees – they use the odor of flowers to identify a good floral source and forage on it exclusively (37, 38). Floral scents are intrinsically variable: substantial variation in the ratios of the compounds exist even for flowers produced by the same plant (39). A honeybee can learn the odor of a flower and subsequently compares it to odors emanating from novel flowers and decides whether or not to forage. It uses several aspects of the odor, (i.e., types of compounds and their concentrations) to

generalize what is learned to novel flowers (39, 40). In a recent study (41), it was shown that honeybees have the ability to use precise information about the ratios of two odors in a mixture in order to identify a rewarding stimulus and to discriminate it from a punishment-associated stimulus having a different ratio of the same two odors. Alternatively, if two stimuli differing in odor ratios both lead to a rewarding outcome, honeybees can learn to ignore information about the ratio of the two odors and thereafter respond to all mixtures of the same two odors with equal probability (41).




Fig. 1.1. Locust Olfactory Anatomy. AL, antennal lobe; LH, lateral horn; MB, mushroom body; OL, optic lobe; agt, antennal-glomerular tract; an, antennal nerve; gl, glomerulus; on, ocellar nerve; p, pedunculus; β LN, β -lobe neuron; KC, Kenyon cell; LN, local neuron; ORN, olfactory receptor neuron; PN, projection neuron; d, dorsal; l, lateral; mid, midline. Adapted from (42, 43).

1.3 Olfactory Microcircuits

1.3.1 *Olfactory receptor neurons*

Olfaction can be thought of as a series of representations and re-representations of the external olfactory world that ultimately allows the nervous system to use it for perception and behavior. The primary representation of an odor is in the response profiles of olfactory receptor neurons (ORNs, ~1300 ORNs from each antenna project bilaterally to the AL in *Drosophila*). The number of unique olfactory receptor (OR) types is very large in most species (~60–1000), making olfaction fundamentally different to other sensory modalities (13). Most ORNs express only one OR (some 2–3), but the same OR is never expressed by more one ORN type (20). The signaling of specific ORNs thus reflects the activity of specific odor receptors. From systematic studies of ORN responses to many odors (44) three basic principles emerge: (i) individual odors activate subsets of ORNs; (ii) individual ORNs are activated by subsets of odors with varying breadth of tuning. Broadly tuned receptors are most sensitive to structurally similar odors; (iii) increasing odor concentrations elicit activity from greater numbers of ORNs. Thus, both odor identity and intensity are represented combinatorially across the receptor population. However, there also exist very specialized channels that are highly selective to certain chemicals. In moths, there are dedicated to pheromones or plant odors (32); in the female malaria mosquito, one

ORN type is highly selective for 4-methylphenol which is present in our sweat (45); CO₂ as low as 0.1% also activates a single glomerulus (V glomerulus) in flies and in turn induces robust avoidance response, CO₂ are released by flies under stress (46).

1.3.2 *Antennal lobe*

In insects, ORNs send axons, via the antennal nerve (AN), to the antennal lobe (AL), where they contact the principal neurons of the AL, projection neurons (PNs, cholinergic), as well as local neurons (LNs, GABAergic, but, in *Drosophila*, some also cholinergic (47)). The neuropil where these synapses are formed is organized into glomeruli (Fig. 1.1). In most insects, ORNs that express the same receptor converge upon one or two glomeruli (32). The convergence ratio from ORNs is high, ~50 ORNs onto an average of ~3PNs per glomerulus (each ORN contacts all PNs in the glomerulus) and with reliable synapses in *Drosophila* (48). This is thought to increase signal-to-noise ratio in PNs (making one PN more informative than one ORN). Most individual PNs receive direct excitatory input from only one type of ORN, expressing one type of odor receptor, in locusts, each PN receives input from 10–14 glomeruli (out of ~1000) with ill-defined boundaries; correspondingly, individual ORN axons project to several glomeruli (32). It is not known whether all ORNs of the same type converge to the same glomeruli (as in *Drosophila*), nor

whether the glomeruli visited by each PN are innervated by ORNs of the same type. Both PNs and LNs receive direct excitatory synapses from ORNs. There are also neuromodulatory neurons that release neuropeptides such as dopamine, octopamine and serotonin, in the AL; these releases are believed to alter PN responses during associative learning (49).

PNs form direct excitatory synapses onto LNs, LNs in turn inhibit PNs, together they form a recurrent network (32). This interplay between excitation and inhibition gives rise to AL dynamics at two time scales (50, 51). First, PNs respond to odors with slow temporal spike patterns that outlast the odor stimulus (12). These patterns include both excitatory and inhibitory epochs. Second, PNs transiently synchronize with each other during an odor presentation (52). This transient synchrony gives rise to ~20–30Hz oscillations seen in the local field potential (LFP: representative of summated synaptic potentials of the PN outputs measured in the Mushroom body). These odor-evoked oscillations are also visible in the subthreshold activity of PNs, LNs and the recipients of PN output in the MB, the Keyon cells (53). GABA_A conductances in LNs are thought to underlie these oscillations, which can be abolished with application of picrotoxin, a GABA_A-like chloride channel blocker (53).

1.3.3 *Antennal lobe-to-Mushroom body circuit*

PNs send axons out of the AL, via the antennal glomerular tract (AGT), and synapse onto the dendrites of Kenyon cells (KCs) in the calyx of the mushroom body (~50,000 KCs in each locust MB). Beyond the MB, the bifurcating axons of PNs also target the Lateral Horn (LH), where they contact inhibitory neurons (LHIs) (54). The MB is a bilaterally symmetrical structure consisting of a calycal neuropil, which forms a cup beneath a large number of KC somata, and a pedunculus that terminates in two or more lobes. The KCs send their dendrites into the calyx and their axons make up the pedunculus and subsequently bifurcate into the lobes. The input to the calyx is predominantly olfactory, from PNs. The output of the MB appears to be restricted to the lobes, where KC axons contact MB extrinsic neurons. The KC population can be divided into multiple types, as determined by anatomical methods (55). In the pedunculus and lobes, KC axons appear to segregate into multiple concentric or parallel layers, subsets of which are selectively invaded by individual MB extrinsic neurons. The MB has been compared to three different regions in the mammalian brain (56): first, the hippocampus, because of its involvement in learning and memory; lesioning the MB in cockroaches impairs their memory for spatial locations, much like hippocampal lesions do in rodents (57). Second, the cerebellum because of its involvement in learning precisely timed motor movements. And thirdly, it has been

compared to the piriform cortex because both are two synapses downstream of the olfactory sensory layer.

KCs are activated by subsets ($50 \pm 15\%$) of coincident PNs, with high threshold, and integrate over $\sim 1/2$ an oscillation cycle (25 ms), are driven most strongly during the most dynamic epochs of PN firing (58, 59). Furthermore, voltage-gate channels in KC dendrites amplify responses to coincident PN inputs nonlinearly, further narrowing the integration time-window (60). These circuit operations between PNs and KCs results in a marked transformation from broadly tuned cells (PNs) to highly odor-selective ones (KCs) (53, 61). This sparse, selective property of KCs has great benefits for memory storage, because if a neuron (i.e., PN) responds to multiple odors, synaptic plasticity driven by one odor could perturb memories formed by a different odor (a problem known as synaptic interference). Such sparse codes have also been found in many other systems (11, 62, 63), and more recently, odor coding has also been found to be sparse in the pyramidal neurons of the piriform cortex (64), an analogue of the MB in rodents.

In many animals, certain odors elicit innate behavioral responses in addition to pheromonal responses (20). Both insects and mammals can be innately attracted to or repelled by certain odors through a mechanism that depends on the activation of specific glomeruli (20, 65). These responses are mediated by hard-wired circuits that link specific ORNs to

specific neurons in higher centers - (i) PN axons from the same glomerulus show stereotyped projections to the LH (66); (ii) PNs associated with food odors and pheromones target different regions of the LH (67). In comparison, a recent electrophysiological study did not find a high level of stereotypy in the MB (68). Experiments that lesion or inactivate the MB suggest that information flow through the LH alone is sufficient to support basic olfactory behaviors (69), while the MB is required for associative olfactory learning. In *Drosophila*, there are ~ 50 MB extrinsic neurons that decode the MB neuronal output (48), in locusts, there likely more (Stijn Cassenaer, personal communication). This shows that after the fan-out (PNs to KCs), there is now again a fan-in (KCs to MB output) as the system is closer to motor output. These synapses are likely changed during associative learning (48). It was recently discovered that there exists a form of spike-timing-dependent plasticity (STDP) at the synapse between KCs and a class of output neurons, termed β lobe neurons (β LN) (70). Synapses that were active a few milliseconds before these β LN spiking were strongly potentiated, while those synapses active shortly after the spike were depressed.

1.4 Outline and Specific Aims

How does the brain achieve both fine recognition of particular odor objects (selectivity) and generalization across categories of odor objects (invariance)? The focus of this thesis is in answering the computational aspects of the aforementioned question, using the locust olfactory system as a model system. The approach I will take is to investigate how neural representations of odor components and mixtures are transformed from the projection neurons (PNs) of the Antennal lobe to the Kenyon cells, intrinsic cells of the Mushroom body, and how KC population data could be read out. In particular, the emphasis of this work will be on the analysis of population neural data of both PNs and KCs. A PN-KC model is implemented and fed with experimental PN inputs that will demonstrate at a detailed mechanistic level, how KC properties observed experimentally could be derived.

The work presented in Chapter 2 builds on a large body of previous research that has elucidated the roles of different elements of the locust olfactory system and described some of the rules governing their interactions (7, 42, 43, 53, 58-60, 71, 72). This work has led to a fairly detailed understanding of the mechanisms that underlie the integration of ensemble PN input by KCs in the MB. But insofar, no detailed and systematic examination of multi-component mixtures have been undertaken in the locust olfactory system, nor has there been any

experimental study of odor segmentation in any neural system (with mixtures beyond 3 components). I will examine the following questions: What are the neural representations of odors and mixtures in the AL and how are they re-represented again the MB? Can we reproduce the characteristics of KCs that we observe experimentally from what we know about the biological constraints of the system.

Chapter 3 applies for the first time a population decoding method to Kenyon cells. What is crucial are the informational aspects of odor stimuli (e.g., category vs. identity) that are represented in the KC population response, because this representation is used by downstream neurons for associative learning and subsequent behavior output. Chapter 3 is organized to answer four questions: (i) What types of olfactory information is extracted by KCs from the PN input? Or put it another way, what types of information is contained in the KC population response? (ii) What is the format of this representation? And how is it different to the format of the PN representation? (iii) What is this format useful for? (iv) How is the KC population response read out by downstream neurons? The work in Chapter 2 and 3 is in preparation for publication as *Shen, K., Tootoonia, S. and Laurent, G.*

Chapter 4 concludes the thesis and presents some open questions for the future.

CHAPTER 2

Mixture Coding and Odor-segmenting Kenyon Cells

The main computational problems of olfaction include odor discrimination (29–31, 73, 74), concentration-invariant recognition (7, 75, 76), classification (grouping of stimuli by shared features), generalization (assignment of novel stimuli to a group, based on shared features), and odor segmentation (of components from within a mixture, of signal from background) (32, 77). These object recognition problems (2) are not specific to olfaction but they are interesting to study from within it, because olfactory systems are structurally shallow and thus solve them in very few steps. Using locusts as models, we gained some understanding of the representation formats for simple odors in the first three relays of its olfactory system—the antennal lobe (AL), mushroom body (MB) and beta lobe (β L)—and of the computations carried out by these circuits (7, 53, 58, 70). We also discovered that odors at different concentrations generate low-dimensional manifolds of spatio-temporal representations (7), providing a neural substrate for concentration

invariance. In this study, we turn to odor mixtures. Most natural odors comprise many components, usually mixed in particular ratios. Mixtures can be perceived as wholes (“coffee”, “grapefruit”) (33), but they can also be classified into categories, with various degrees of refinement (“fruity” → “citrusy” → “grapefruit”). In addition, humans can identify as many as 8–12 familiar odors in a blend (33), animals such as insects and rodents can likely do better (78). These observations are interesting, because the computational constraints on generating a unitary percept and on segmenting a stimulus into its components are contradictory. Our goal was thus to discover, using the locust system, whether and how the formats of representations for odor mixtures might be consistent with these competing requirements. In the first set of experiments, we investigated neural representations to two odors and analyzed how these responses changed when the stimulus was “morphed” from one odor to another through a series of intermediate mixtures (binary mixture experiments, see Methods). In the second set of experiments, we chose a set of eight monomolecular odors, paraffin oil (their dilution substrate) and mixtures of two, three, four, five and eight of those odors (44 stimuli in all, out of 211 possible, see Methods). We recorded from 343 projection neurons (PNs, the analog of vertebrate mitral cells, 168 PNs for binary experiments, 175 PNs for multi-component mixture experiments) and 209 Kenyon cells (KCs, the mushroom body neurons, for multi-component mixture experiments) in 61 animals.

2.1 Results

2.1.1 *Representations of binary mixtures by PNs*

Because odor representations by PNs are highly distributed and varied in time (Fig. S2.3), because their activity patterns are decoded by individual KCs on which converge many PNs (59) and because KCs have very short effective temporal integration windows (53, 60), it is appropriate and more informative to examine PN responses as time-series of instantaneous population vectors, or trajectories, in an appropriately reduced state space (7, 58, 71, 79). Figure 2.1 illustrates these PN population trajectories for a set of representative stimuli. Figures 2.1A–C concern binary mixtures: we plot the evolution of the representation for the odors citral, octanol, and their 1:1 mixture. The mixture trajectory lies somewhere in between those for the two components, suggesting a simple linear combination. This was confirmed by correlation analysis performed in full PN space. This relatively simple combination was not entirely predictable from the responses of single PNs to binary mixtures, for those often deviated significantly from the arithmetic sum of the responses to the components (Fig. S2.3; compare open and filled PSTHs). This suggests that significant correlations exist between the responses of different PNs to the same stimulus, and that those linearize the population's combined output, at least for binary mixtures. Figure 2.1B represents concentration series for these three stimuli (the two pure

odors and their binary mixture). Extending previous results (7), we find that concentration series for 1:1 mixtures, as for single odors, generate families of closely related trajectories (lower-dimensional manifolds), clustered by odor rather than concentration. In a final experiment, we “morphed” one odor into the other in 11 intermediate steps (Fig 2.1C). Contrary to recent results in the zebrafish olfactory bulb (Friedrich et al., *in press*), we observed no sudden transition but rather, a gradual shift of the population trajectory corresponding to one odor to that for the other odor, via their 1:1 mixture trajectory. Thus the encoding space defined by PNs appears to optimize the spread of odor representations to accommodate even small changes in the stimulus. While the responses of single PNs often deviate from the linear combination of the responses to their components, the population output is reasonably well approximated by linear summation.

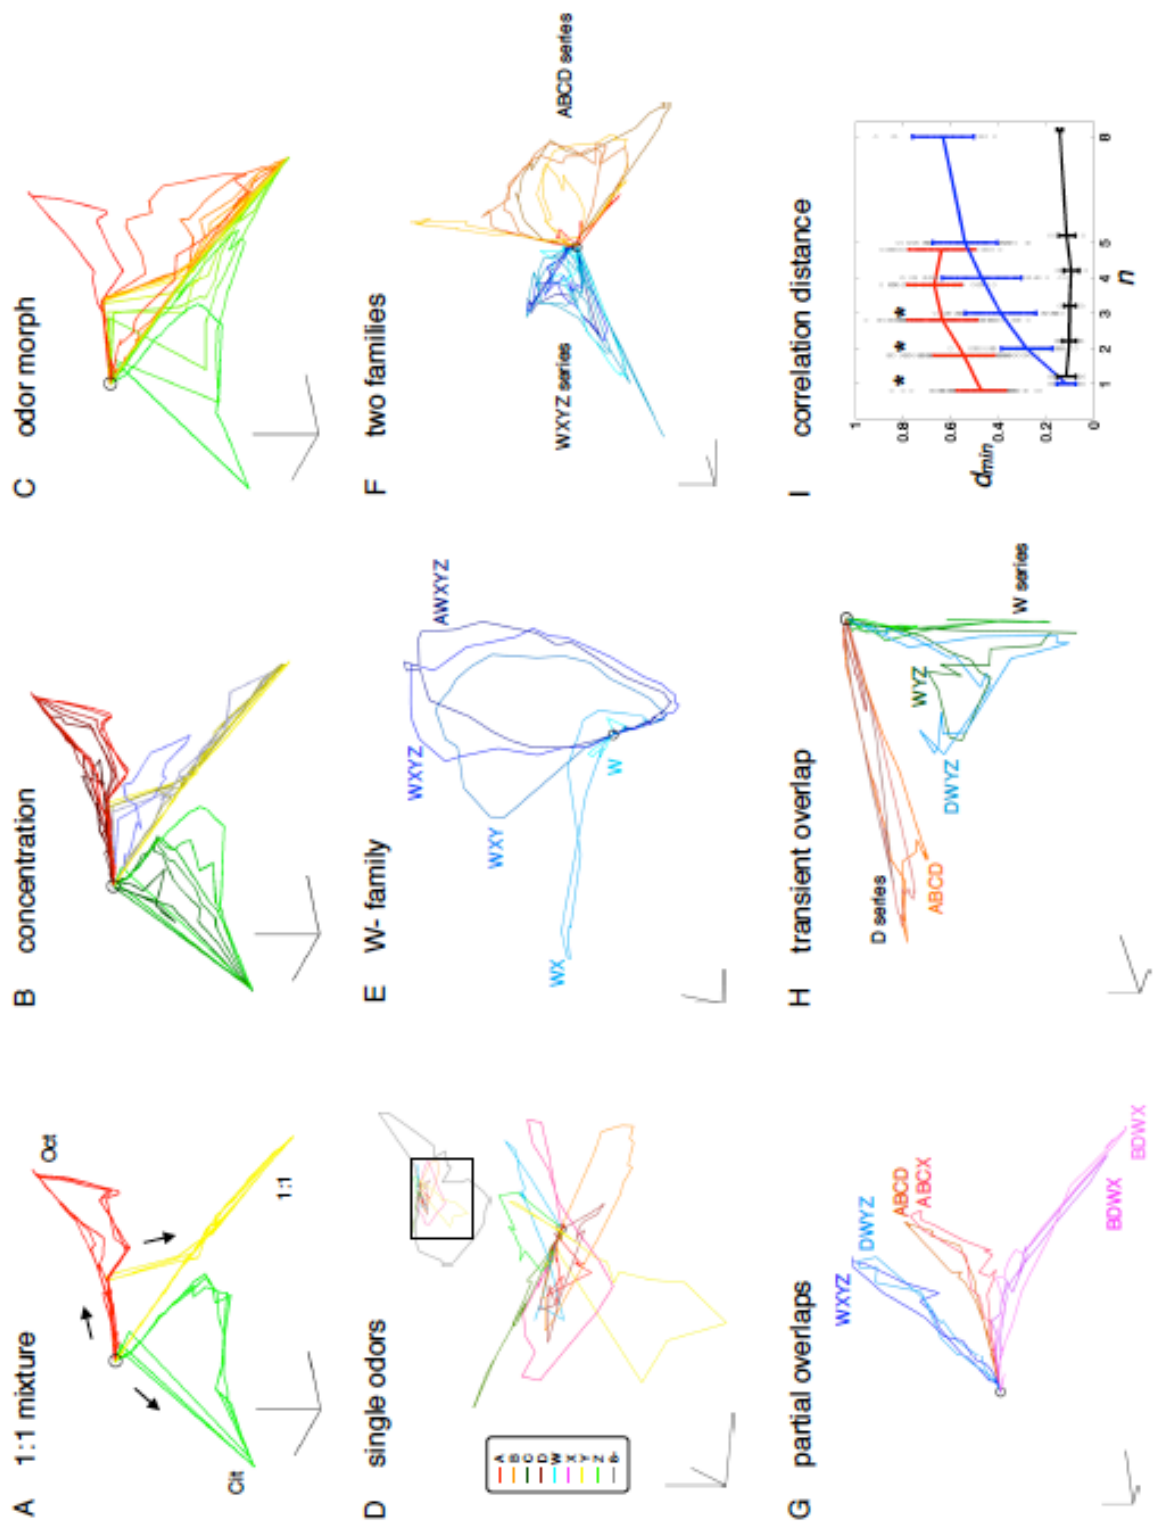
Figure 2.1 Representations of single odors and their mixtures are spread orderly in PN coding space.

A–C. Trajectories representing PN-ensemble responses to binary mixtures. PN activity is represented as a point in 168-D space, where each dimension represents the firing rate of one of the 168 PNs during one 50-ms time bin. Data analyzed using LLE and projected in the space of the first three LLE components (see Methods). Arrows indicate direction of motion. Three seconds are represented, beginning at odor onset; odor pulses are 300 ms long; each trajectory composed of a sequence of 50ms-bin measurements, averaged over three trials. **(A)** PN Population responses to single odors (citral: green; octanol: red) and to their 1:1 mixture (yellow). Initially at a resting state (origin: O), the PN population responds with stimulus-specific trajectories. **(B)**

Trajectories to different concentration series of citral, octanol (30, 60, 80, 100, 120, 140 ml/min) and their 1:1 mixture (30:30, 60:60, 80:80, 100:100, 140:140 ml/min). Similar to trajectories for pure odors, concentration-specific trajectories of the 1:1 mixture form odor-specific manifold. Nine-trial averages for each condition. **(C)** Trajectories corresponding to odor-morphing series. From Oct to Cit: 140:0, 140:30, 140:60, 140:80, 140:100, 140:120, 140:140, 120:140, 100:140, 80:140, 60:140, 30:140, 0:140. Trajectories change smoothly, with greatest changes away from pure-odor.

D-H. Multi-component-mixture trajectories (dataset different from that in A-C: 175 other PNs, stimulated with 44 different odor conditions, see Methods). Four and a half seconds are represented, beginning at odor onset; odor pulses: 500 ms long; 50-ms bins, each averaged over three 2-trial averages. **(D)** Trajectories in 3-LLE space for the 8 single-odor components. Inset: zoom out of 8 single-odor trajectories together with the 8-mixture trajectory (gray) (LLE axes recalculated). Mixture trajectory loops around those for individual odors. **(E)** Starting from single odor W, trajectories increasingly deviate as components added ($W \rightarrow WX \rightarrow WXY \rightarrow WXYZ \rightarrow AWXYZ$). **(F)** Mixtures form ordered trajectory clusters: family of {W,X,Y,Z} (W, X, Y, Z, WX, WY, WZ, XY, XZ, YZ, WXY, WYZ, WXYZ) well separated from family of {A,B,C,D} (A, B, C, D, AB, AC, AD, BC, ABC, ACD, ABCD). **(G)** Trajectories to partly overlapping 4-mixtures. **(H)** Four D-containing trajectories (D, AD, ACD, ABCD) plotted together with four Z-containing trajectories (Z, WZ, WYZ, DWYZ). $\underline{D}WYZ$ trajectory (cyan) follows WYZ trajectory for the most part but deviates towards the D-series transiently. Hence, DWYZ can be classified as related to Z or D, depending on time within response (see text).

I. Minimum correlation distances (d_{min}) between two trajectories as functions of number of components in the mixture (correlation distances minimize contributions of firing rate differences). Each bin of a trajectory (2-trial average) is compared with bins of the other at times $t \pm 2\text{bins}$. Correlation distances in space defined by the 40 first principal components ($\sim 70\%$ of variance). Details in text. When comparing distances between single odors and mixtures, upper bound is given by distances between different single odors (red, $n=1$), because this distribution indicates the separation between non-overlapping stimuli.



2.1.2 *Representation of mixtures of increasing complexity by PNs*

We next examined PN trajectories for mixtures of increasing number of components. Eight molecules were chosen to be chemically distinct and their concentrations adjusted to evoke minimal, reliable and comparable electro-antennograms, compensating for differences in vapor pressure or receptor activation and ensuring operation away from saturation. The trajectories corresponding to these eight stimuli are shown in Fig. 2.1D. Consistent with the odors' distinct chemical composition, these trajectories did not cluster, indicating large differences between the evoked PN response patterns.

We first examine the effect of adding $1 < n < 7$ components to a single odor, W (Fig. 2.1E). The mixture trajectories always deviated from that for W and from each other. For $n > 3$, however, subsequent component addition led to decreasing changes in the population trajectory. This is consistent with the fact that the fractional change to the stimulus decreased with each single component addition. This observation was repeated with the other odors and quantified by analysis in the full PN space (not shown).

Second, we observed that, while mixture representations deviated from those of their components, they still formed clusters of trajectories, well segregated from those corresponding to non-overlapping mixtures. In Fig. 2.1F, sets of all single- and mixed-odor trajectories for odor groups

{W,X,Y,Z} and {A,B,C,D} are plotted, revealing two non-overlapping manifolds. This suggests that PN population patterns retain information about components in mixtures, and that PN trajectories do not spread randomly in representation space.

Third, we examined the trajectories corresponding to partly overlapping odors with equivalent strengths (same numbers of components). In this example (Fig. 2.1G) we plot the trajectories of six mixtures. Three pairs had an overlap of 3 out of 4 components (BCWX & BDWX; ABCD & ABCX; WXYZ & DWYZ) and these pairs clearly clustered together. The other combinations overlapped by two (e.g. BCWX vs. ABCD; BCWX vs. WXYZ) and were roughly equidistant from one another. This again suggests an ordered occupancy (qualitatively at least) of PN coding space, where distances between population representations decrease as composition overlap increases. Note, however, that overlaps between mixtures representations—considered until now as averages—often changed over the course of a trajectory. Figure 2.1H, for example, plots the trajectories for two groups of odors that were distinct (ABCD vs. WYZ), until component D was added to WYZ. The addition of D caused a new kink in the DWYZ trajectory, bringing it closer to the D family during a short segment of the response. Conversely, two highly overlapping mixtures (overlap of 4 out of 5 components: ADWYZ and AWXYZ) could be represented by PN trajectories that remained nearly identical for a segment of the response,

and then split apart over a later epoch (Fig. S2.8). These results showed that a fair metric of the similarity between two mixture representations should be based not on the totality of their corresponding trajectories, but on piecewise measurements, and on the closest encounter between them. We thus measured the minimum correlation distances between every single-component odor and all other odors (other singles, mixtures of 2, 3, etc.). Using correlation distance (as opposed to Euclidean) has the advantage of focusing on differences in PN population vectors and discounting effects attributed to changes in firing rate (such as concentration). This minimum-distance plot (Fig. 2.11) was calculated in three ways: between trials (black), to measure the variance of individual population representations; between the representations of each single component and those of all the mixtures containing it (blue); between the representation of each single component and those of all the mixtures excluding it (red). The blue and red curves (and corresponding distributions) were significantly different (Wilcoxon Rank Sum Test, $p < 0.01$) only for $n=1, 2$ or 3 . Thus, the representations of a monomolecular odor and of mixtures of $n>3$ components are equally distant (on average over odors, and at the times corresponding to minimum distances) whether the mixture contains that component or not. In conclusion, while PN-representation space clearly shows order from mixture coding, extraction of component composition, based on overlaps between PN population vectors appears difficult if not impossible with mixtures of $n>3$ components.

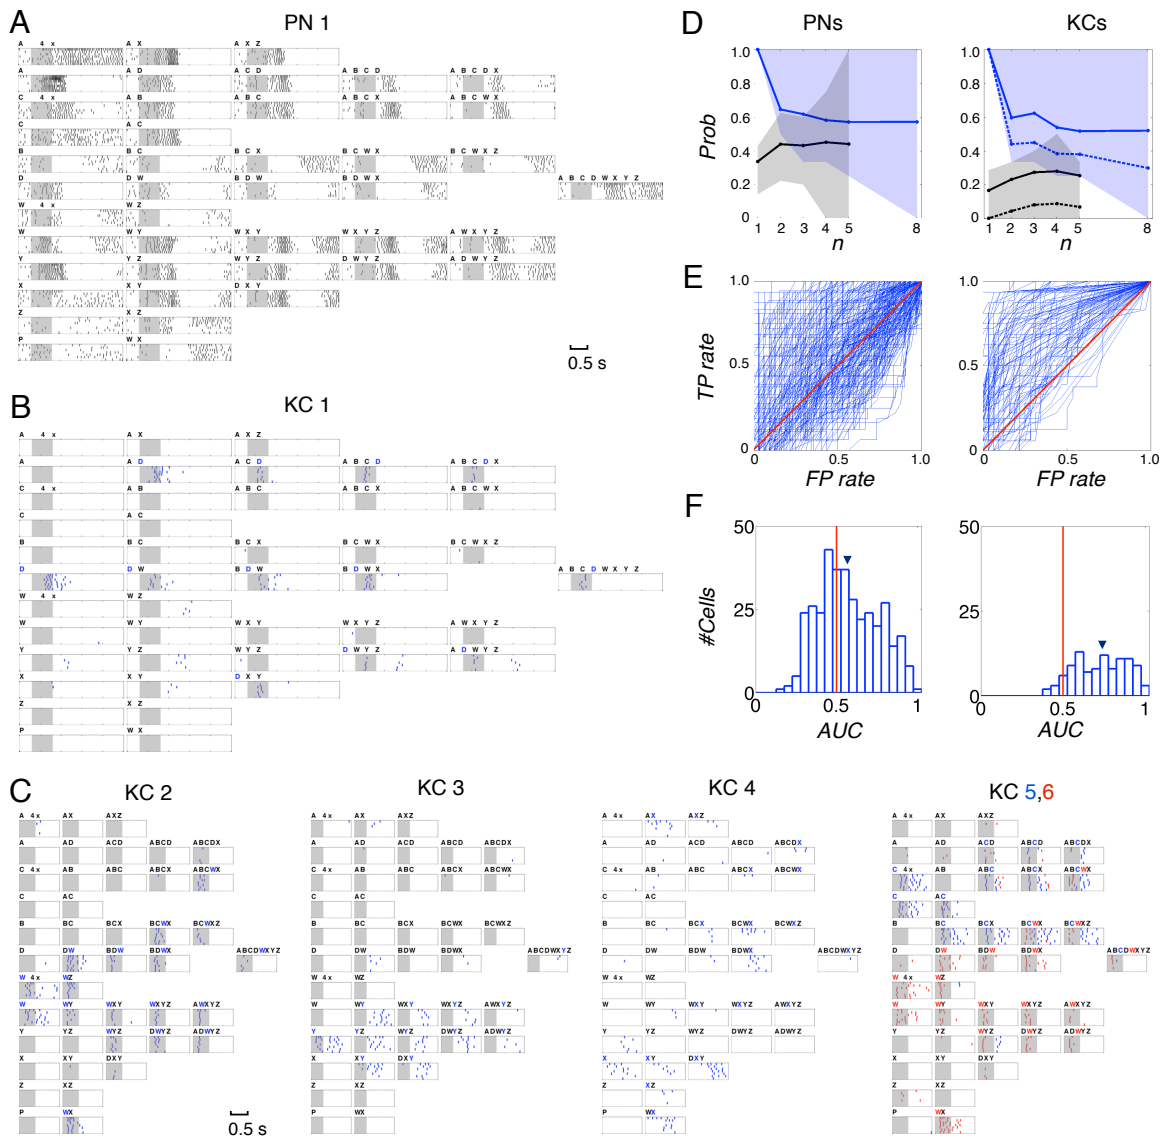


Figure 2.2 KCs segment components out of odor mixtures, but PNs do not.

A. Spike rasters of a representative PN to single and mixed odors (see Methods for 44 stimuli, 7 trials, 500–ms stimulus at shaded area, 2.5 s shown). Numbers of components organized by column, conditions arranged so that overlapping mixtures next to each other wherever possible. Increased inhibition often observed for mixtures

with increasing size; responses to mixtures difficult to predict from responses to components (see Fig. S9).

B–C. Spike rasters of six representative KCs (see Figs. S11–14 for more examples). **(B)** D-segmenting KC, with weak late response to unrelated mixtures Y, WZ, YZ. Same scale as in A. **(C)** Other representative KCs, showing segmentation of different odors (in order of KCs: W, Y, X, C and W). KCs 5 and 6 recorded simultaneously; both responded to mixtures containing both C and W (e.g., BCWX), but at different times. Only 1 s shown, centered on KC response times; t scale as in B.

D. Conditional probability of response to mixtures, given that cell responds to components (see text and Method). Blue: “in-class”. Black: “out-class”. Lines: “inclusive”; dashed: “exclusive”. Averaged across all responding cell-odor class pairs. Shaded region: 30–70% distribution. Separation between in-class and out-class much greater for KCs than PNs.

E. ROC evaluation of component selectivity by PNs and KCs. True and false positive rates (TP, FP) determined by sliding response threshold; response based on spike counts within 1 s window summed over 7 trials. Red diagonal: chance performance. Blue lines: results for each responding cell-odor class pair (see Methods for class partitions).

F. Distribution of area-under-curve (AUC) values for KC-odor class pairs significantly shifted to the right of PN-odor class pairs ($p < 3 \cdot 10^{-13}$, Wilcoxon rank sum test). Arrows indicate means: 0.73 (KCs; SD=0.16); 0.57 (PNs; SD=0.18).

2.1.3 *Kenyon cell responses to mixtures*

Because Kenyon cells are the direct targets of PNs in the mushroom bodies, because mushroom bodies are a site for associative memory (48, 80) and because KC output synapses are plastic (70), KCs are the likely repository of olfactory memories. It is therefore important to determine the stimulus features that they extract from PNs. For comparison, we show first the responses of one representative PN to our 44 stimuli. As is typical of PNs

(53), this neuron responded to about half of the stimuli with a variety of discharge patterns (Fig. 2.2A, see also Fig. S2.10). By contrast, KCs responded very rarely to single odors, but when they did, did so with very high specificity (Figs. 2.2B–C). Surprisingly, KCs that responded to a component also often responded to many—if not all—of the mixtures containing it. KC 1 (Fig. 2.2B), for example, fired in response to odor D, and responded to all mixtures containing D (though not necessarily at the same times and for the same durations). The same can be seen with KCs 2 and 6 for odor W (Fig. 2C). KCs 5 and 6 were recorded simultaneously, and each responded to a different molecule. We found KCs specific to all 8 single odors. (Our pre-experiment search for KCs always focused on these 8 odors, but on them only, see Methods; we also found, by chance and thus rarely, a few KCs specific for binary mixtures; Fig. S2.13). Thus the ability to detect components in a mixture appears to occur first with KCs.

We next analyzed the difference between PN and KC responses using two metrics. In the first, we measured conditional probabilities of response to mixtures, given that a neuron responded to a component (see Methods for definition of response). If a neuron responded to component c , we measured the fraction of c -containing mixtures that it responded to (blue curves). This was repeated for all component–cell combinations with PNs and KCs (Fig. 2.2D). With KCs, this measurements were performed in two ways: an “inclusive” computation contained all cells responding to *at least*

one component (continuous lines); an “exclusive” computation contained KCs that responded to *only* one component (stippled line); the later measure is more informative since it excludes the potential contribution of components other than that tested on responses to mixtures. This exclusive computation was not possible with PNs for they always responded to many components. Our second approach was a receiver-operator-characteristics (ROC) analysis (81), measuring a neuron’s ability to separate stimuli into “containing- x ” and “not-containing- x ” sets, as response threshold is varied. On a true-positive (TP) vs. false-positive (FP) plot, selective neurons are identified by ROC curves located in the upper-left quadrant (Fig. 2.2E). Unselective ones run along the diagonal. The area under the curve (AUC) thus measures selectivity (near 1 for high, near 0.5 for low) (Fig. 2.2F). Both approaches indicated that KCs are significantly better than PNs at component segmentation. ROC analysis proved that this is not explained simply by high KC firing thresholds. Hence, in addition to being highly selective and thus, rare responders, KCs behave as odor segmenters, extracting component information from PN population vectors.

2.1.4 PN and KC population statistics

We quantified population PN and KC activity as a function of n number of odor components in the mixture (Fig. S2.17A). Mean baseline PN firing rate calculated was ~ 2.5 Hz. For single components, peak firing reached ~ 3 Hz,

around odor offset. Peak instantaneous firing (and total spikes) remained approximately the same as a function n , interestingly, peak firing rate increased for the higher concentration of single components (~ 3.8 Hz), while mixtures of comparable concentration resulted in a lower firing rate (~ 2.8 Hz for 4-mixtures, Fig. S2.17A). A closer examination of the response profiles of firing rates reveals that the onset of firing is earlier for mixtures than for components, and this difference cannot be accounted for by concentration alone, e.g., compare peak and onset of 3- mixtures to 4x 1-components. Next, we examined the percentage of silent PNs as a function of time and n components. A cell was defined as silent if it fired no spikes in 100 ms time bins (across 7 trials) to allow for a more conservative measure of silence. The percentage of silent PNs clearly increases as a function of n components (Fig. S2.17B). This reflect increased inhibition by local neurons (LNs) onto PNs. When many components are mixed, this inhibition is greater for mixtures than for components of comparable concentrations. Together, these results suggest a gain control mechanism of mixtures mediated by the PN-LN network that regulates the output of the PNs. Next, we examined the firing rate of the KC population as a function of n (Fig. S2.17A), unlike PNs, instantaneous KC firing increases as a function of n , for ~ 0.3 Hz for single components to ~ 0.9 Hz for 8-mixtures. In addition, we observe that unlike PNs, where many cells become inhibited during odor response, most KCs by comparison are silent at rest, and a very small percentage of them become active during response. This small, but

still significant increase in KC firing as a function of n must be attributed to not greater number of PN spikes, but greater synchrony of PN inputs.

An important property of neural codes is the activity ratio, the fraction of active neurons at any one time (82). At one end of the spectrum are local codes, where each stimulus is represented by a single active cell. At the other extreme are dense distributed codes, where each stimulus is represented on average by about half of the cells, e.g., the ASCII code. Codes with low activity ratios are known as sparse codes, where each stimulus is represented by much a smaller neuronal population (but not 1 cell), the members of which respond in an explicit manner to specific features. The activity ratio has implications for coding capacity, memory recall, generalization, fault tolerance and speed and rules of learning (see Ch 3). Here, we compared the activity ratio between PNs and KCs (Fig. S2.17C). We measured the responsiveness of cells in short time bins of 50 ms. A cell was defined as responding if it spiked at least once in 4 of 7 trials, and was at least 1.5 SDs above the baseline firing (see Methods). At rest, less than 0.4% of all PNs are responding by this metric, however, with odor onset, the percentage of responsive PNs immediately rose to ~8% for single components, ~13% for higher concentrations of single components, and ~11% for multi-component mixtures. Because the identities of responding PNs change from time bin to time bin, over a 3 s period, ~36–55% of all PNs

responded in at least one time bin (this percentage increased as a function of n). In comparison, there were no responding KCs at baseline, and only a very small subset of them responded with odor onset, $\sim 0.5\text{--}1\%$ in any one time bin, and a maximum of $\sim 5\%$ in a 3 s period (to 8- mixture). These KC response probabilities are qualitatively a 10-, 20- fold over-estimation, because there was a selective bias for single components in the KC recordings (see Methods). These results confirm that PN responses are dense and distributed with a mechanism for gain control of PN outputs. In contrast, KC responses are sparse and rarely respond (this issue is revisited in Ch 3). Interestingly, in a recent modeling study (83), it was shown that optimal discrimination performance is associated with a narrow range of values for sparseness centered at $\sim 1\%$, matching our KC responsiveness.

Figure 2.3 KCs segment components out of odor mixtures, but PNs do not.

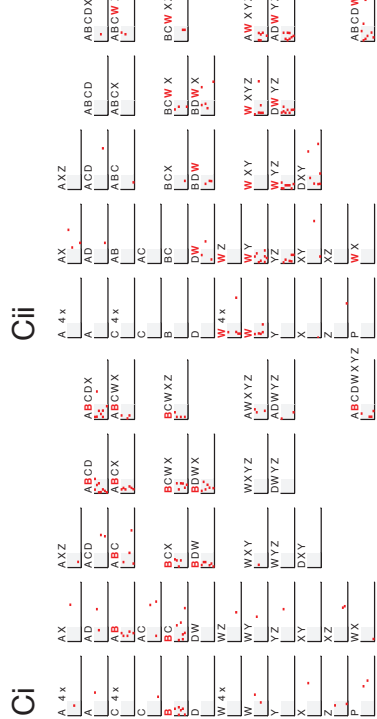




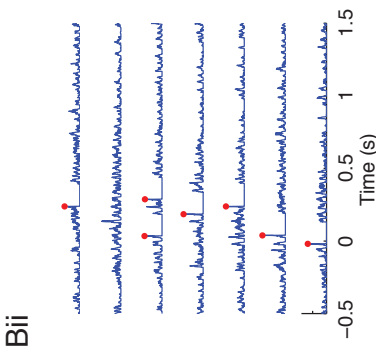
A. Spike rasters of a representative PN to single and mixed odors (see Methods for 44 stimuli, 7 trials, 500-ms stimulus at shaded area, 2.5 s shown). Numbers of components organized by column, conditions arranged so that overlapping mixtures next to each other wherever possible. Increased inhibition often observed for mixtures with increasing size; responses to mixtures difficult to predict from responses to components (see Fig. S9).

B–C. Spike rasters of six representative KCs (see Figs. S11–14 for more examples). **(B)** D-segmenting KC, with weak late response to unrelated mixtures Y, WZ, YZ. Same scale as in A. **(C)** Other representative KCs, showing segmentation of different odors (in order of KCs: W, Y, X, C and W). KCs 5 and 6 recorded simultaneously; both responded to mixtures containing both C and W (e.g., BCWX), but at different times. Only 1 s shown, centered on KC response times; t scale as in B.

D. Conditional probability of response to mixtures, given that cell responds to components (see text and Method). Blue: “in-class”. Black: “out-class”. Lines: “inclusive”; dashed: “exclusive”. Averaged across all responding cell-odor class pairs. Shaded region: 30–70% distribution. Separation between in-class and out-class much greater for KCs than PNs.

E. ROC evaluation of component selectivity by PNs and KCs. True and false positive rates (TP, FP) determined by sliding response threshold; response based on spike counts within 1 s window summed over 7 trials. Red diagonal: chance performance. Blue lines: results for each responding cell-odor class pair (see Methods for class partitions).

F. Distribution of area-under-curve (AUC) values for KC-odor class pairs significantly shifted to the right of PN-odor class pairs ($p < 3.10^{-13}$, Wilcoxon rank sum test). Arrows indicate means: 0.73 (KCs; SD=0.16); 0.57 (PNs; SD=0.18).



2.1.5 *Generating model KC classifiers using recorded*

PN data

A simple abstraction to explain our observations is that odor representations are spread orderly in a high-dimensional PN space (Fig 2.1A); because of their partial and specific connectivity to PNs ($50 \pm 15\%$, (59)), individual KCs sample different lower-dimensional projections of PN space. At every oscillation cycle, each KC makes a binary classification decision in the subspace that it sees, on the presence or absence of a particular stimulus feature (odor component). This abstraction might explain the fact that a single KC can recognize a component, even when the PN trajectories corresponding to the mixtures containing it differ significantly from one another in the full PN space (Fig. 2.1A): by projecting those PN mixture trajectories into the appropriate subspace (by sampling the appropriate subset of PNs), a KC could detect the appropriate “crossings” of the projected trajectories. To test this intuition, we generated a simplified model of the antennal lobe-mushroom body circuits, fed into it our recorded PN data, and tested whether it was sufficient to produce model KC (mKC) responses similar to those recorded in our experiments. Our constraints on the model’s design were entirely determined by our knowledge of the system (53, 59, 60) and by our PN data (see Methods).

There are four main components to our simplified PN–KC model. First, each mKC received direct PN input from $50 \pm 15\%$ of the incoming PNs (59). The large number of different PN combinations implies that there are a large number of patterns that KCs could potentially encode (59). Such convergence (~ 400 PNs–1 KC) and divergence (one PN–many KCs) would lead to a lot of KC firing unless the KC’s threshold were appropriately high and appropriate gain control (see below). Second, EPSPs were modeled as first order filters with amplitude A_1 , and decay τ_1 , both subjected to a nonlinearity constrained by electrophysiological results (53, 60). This was to mimic voltage–dependent conductances that serve to sharpen EPSPs when KCs were depolarized. Thus, a PN spike that arrives synchronously with others can contribute disproportionately towards the KC reaching threshold. Third, KCs receive odor–evoked feedforward inhibition from the LHI neurons. These GABA–ergic neurons which respond non–specifically to odors appear to have extensive axonal arborizations in the MB. Each mKC thus received delayed feed–forward inhibitory input from the entire PN population, modeled as an IPSP (A_2, τ_2) with delay Δt , representing these LHIs (53). Overall, KCs receive both excitatory and inhibitory inputs at specific times that are locked to the LFP. PNs tend to fire preferentially during the rising phase of each cycle. LHIs, which receive their input from PNs, tend to fire with a delay and during the falling phase of the cycle. Thus, in each cycle during an odor

response, a KC receives excitatory input from the set of PNs it is connected to that happen to be active during that cycle. Immediately after this, the KC receives non-specific inhibitory input from the LHIs that resets its membrane potential limiting integration of EPSPs across oscillation cycles. Fourth, The respective excitatory and inhibitory inputs were summed and compared to the firing threshold, itself regulated adaptively by the entire mKC population output, through the normalizing negative feedback pathway recently identified (giant GABAergic neuron, Papadopoulou et al.). This feedback gain control (the final component of our model) was indispensable to generated mKC-population response statistics commensurate with experimental results.

Each KC differed from the others only by its connection vector to the PN population, drawn randomly and independently. This model system was fed the PN spike data (phase warped to simultaneously recorded LFP) recorded experimentally with our 44 stimuli. Because our mixture dataset comprised ~151 PNs (vs. 800 in the entire population), the AL-MB model was scaled down to 10,000 KCs (rather than 50,000). The simulations were run 100 times, generating 1,000,000 model KCs (mKC), each uniquely determined by 55 randomly selected PNs out of 151 possible PN inputs to our model. We identified each mKC's response profile and examined the response statistics of this population. The responses of one mKC to a single odor (7 trials) are shown in Fig. 2.3C. The advantage of this approach is that

mKC responses were entirely constrained by PN activity (given by recordings), by our knowledge of the circuits (given) and by PN–mKC connectivity (the only variable). Hence, the null hypothesis was that random connectivity between PNs and KCs is a sufficient constraint to generate component–detecting KCs, and to produce them with the distribution observed experimentally.

Two examples of classifying mKCs are shown in Fig. 2.3C. As observed experimentally, these two mKCs responded to one odor (B or W), and to most mixtures containing the component. Over the population of mKCs, firing rates varied with mixture composition precisely as observed experimentally. Over 10^6 mKCs, however, only 1,200 were found to be segmenters. (Over 50,000 true KCs, this fraction would be equivalent to 60 segmenting true KCs, a number grossly inconsistent with our experimental discovery over a very small sample.) Similarly, there was no separation between conditional response probabilities to mixtures containing the component and those excluding it, contrary to experimental results. Hence, mKCs did not classify odors with the frequency observed in our experiments.

Because the only features distinguishing the responses of mKCs were their connectivity to PNs, we extracted the 1,200 good segmenters from among all mKCs and identified the PNs to which they were connected. For each odor and classifying mKC, we rank ordered the PNs according to

whether they tended to be present or absent in the input vector. With this knowledge, we returned to the full mKC population and imposed a selected bias on the composition of the input vectors to mKCs, thus deviating from randomness. We increased this bias until conditional response probabilities for mKCs matched those observed experimentally. We observed that when 20% of all mKCs connectivity were biased, by manipulating 15/55 of their PN inputs it was sufficient to explain our experimental observations. Hence, a circuit constrained by data, with $50 \pm 15\%$ input connectivity and a small bias away from randomness can fully account for our experimental results on odor classification by KCs.

2.2 Discussion

2.2.1 *Functional consequences*

While the representations of odor mixtures by PN assemblies show clustering by chemical composition, the relationship between the representation of a mixture and that of one of its components is on average no tighter than that between mixture and unrelated components as soon as the mixture contains more than 3 components. Surprisingly KCs—directly postsynaptic to PNs—are individually much better than PNs at detecting a component in a mixture of up to eight odors (ROC analysis). By building a reduced model of the locust PN–KC network constrained by experimental

data (53, 59, 60) and feeding this model our experimental PN data, we showed that the segmenting properties of KCs can be entirely explained, qualitatively and statistically, provided that connectivity between the PN and KC populations, set at $50 \pm 15\%$ by experiments (59), is not entirely random. This suggests either a genetic encoding of PN–KC connectivity or more likely, the existence of a learning rule, presumably unsupervised and yet to be discovered there, to fine-tune PN–KC connectivity in the mushroom body. This model also suggests that we now have a relatively good mechanistic understanding of these early olfactory circuits. Among its key components is an all-to-all normalizing feedback loop within the mushroom body (Papadopoulou et al). This normalizing feedback loop is mediated by a giant GABAergic neuron (GGN) that has extensive arborizations in the MB (Papadopoulou et al). This neuron was found to be non-spiking and provides increased inhibition as the number of components was increased in the odor mixture. In our PN–KC model, we found that without feedback gain control from the GGN, KC firing increased at a much greater rate as a function of n components than what experimental KCs. Because this was the one variable in our model that KCs were the most sensitive to, we conclude that GGN feedback onto KCs is crucial in maintaining sparsity in KCs, and in turn generating component-selective KCs.

2.2.2 *KCs as independent feature detectors?*

Jortner et al. (59) postulated that because the connectivity between PNs and KCs are dense ($50 \pm 15\%$), and if each KC sampled m random PNs out of n PNs, then the number K of possible PN combinations would be maximal, estimated to be $K \approx 10^{240}$, an enormous number. And because there are only 50,000 KCs and assuming each KC samples PNs randomly, then the chance of hitting the same KC would be almost zero, 10^{-240} . On average we would expect 50% overlap between inputs to any two KCs, and this would maximize KC selectivity. Given the experimental findings presented in this thesis, PNs and KCs are most probably not connected in a random fashion. The fact that KCs with similar tuning properties (although limited in the space of 44 odor stimuli) could be found across multiple animals already suggests that KCs might not sample PNs randomly. PNs as a population encode different features of the stimulus in different oscillation cycles shown as trajectories in PN coding space (Fig. 2.1). In every oscillation cycle, KCs sample a subset of PN inputs and makes a binary classification in that space. This results in KCs that are sparse, invariant and explicit representations of stimulus features. One can ask why are so many KCs component detectors?

From a machine learning perspective, given the enormous numbers of features at one's disposal, there is a necessity to reduce the number of features in order to one, reduce computational complexity and two, to

maximize classification performance (19). It has been shown that the higher the ratio of the number of training patterns N to the number of free classifier parameters, the better the generalization properties of the resulting classifier (19). Thus, for a limited number of training patterns (natural encounters of odors), keeping the number of features as small as possible would be advantageous. Given the very large space of all possible features (and corresponding large PN space), are there general rules that KCs use to select the most important features so as to minimize the number of KCs needed, and conserve their ability to best classify and identify different odors? If KCs select features with little discriminatory power, subsequent classifier would be poor. If information-rich features can be selected, the design of the subsequent classifier could be simplified. Feature detection should be computed such that it is as independent as possible. There is no reason for features to be redundant, since it does not improve classifier performance. One possible solution to this problem is to select independent features that have large between-class distances, – one theoretical predication would be that KCs select features that are as independent as possible. The eight odor components used in the experiments discussed in this thesis came from functionally very different groups and I found KCs that responded to all of them (Fig. S2.14–5). These results might allow us to speculate that KCs are independent feature detectors, and groups of KCs form feature vectors.

Fig. S2.1


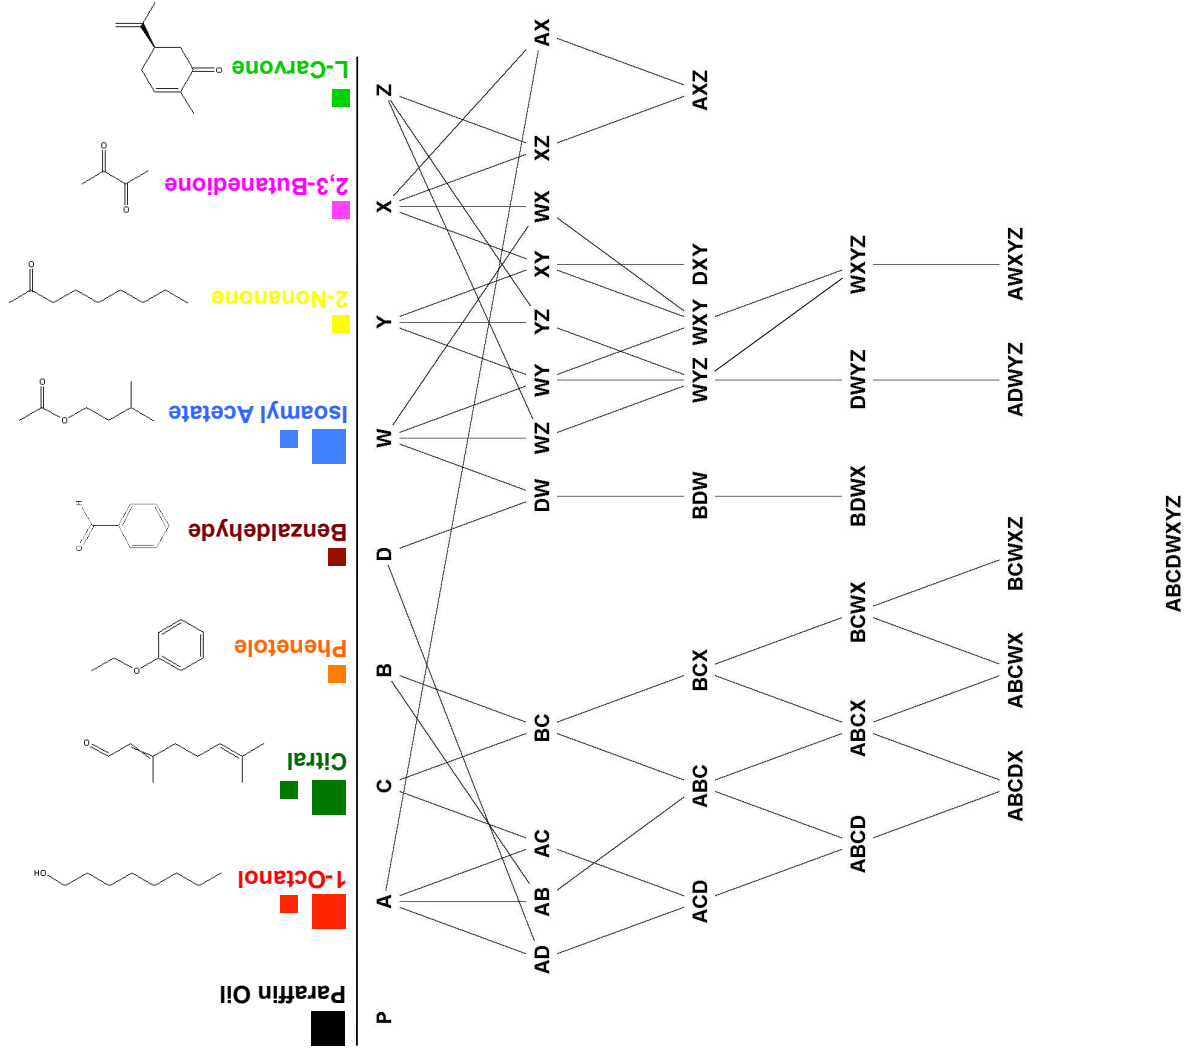


Fig. S2.2



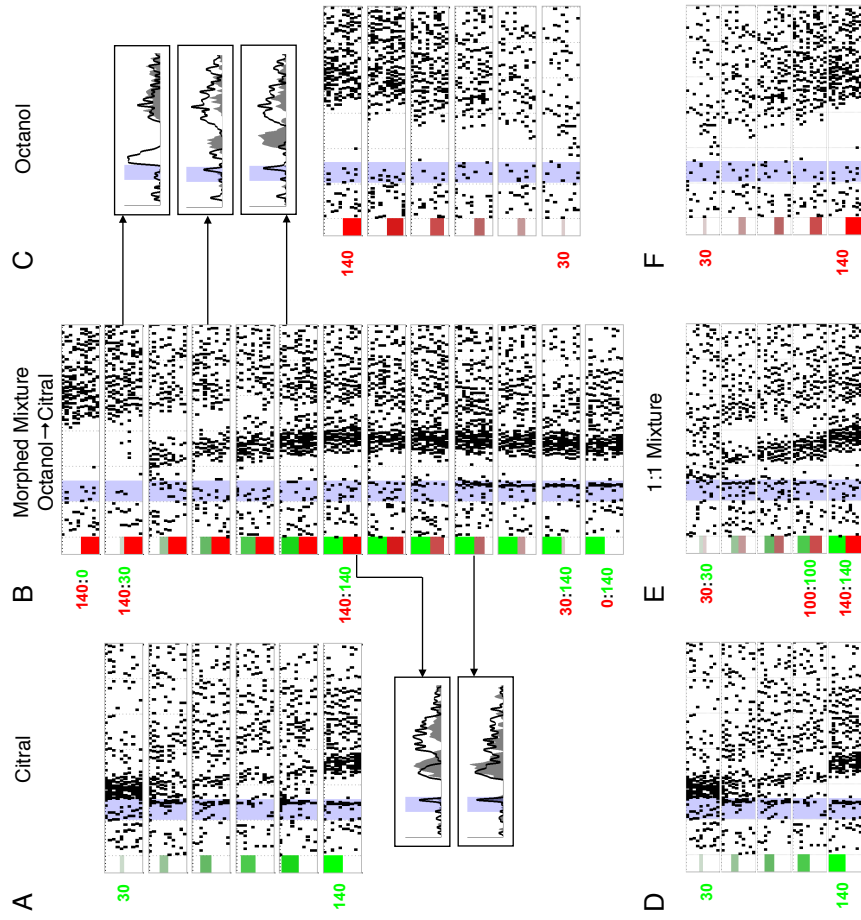



Fig. S2.3

Fig. S2.4






Fig. S2.5




Fig. S2.6

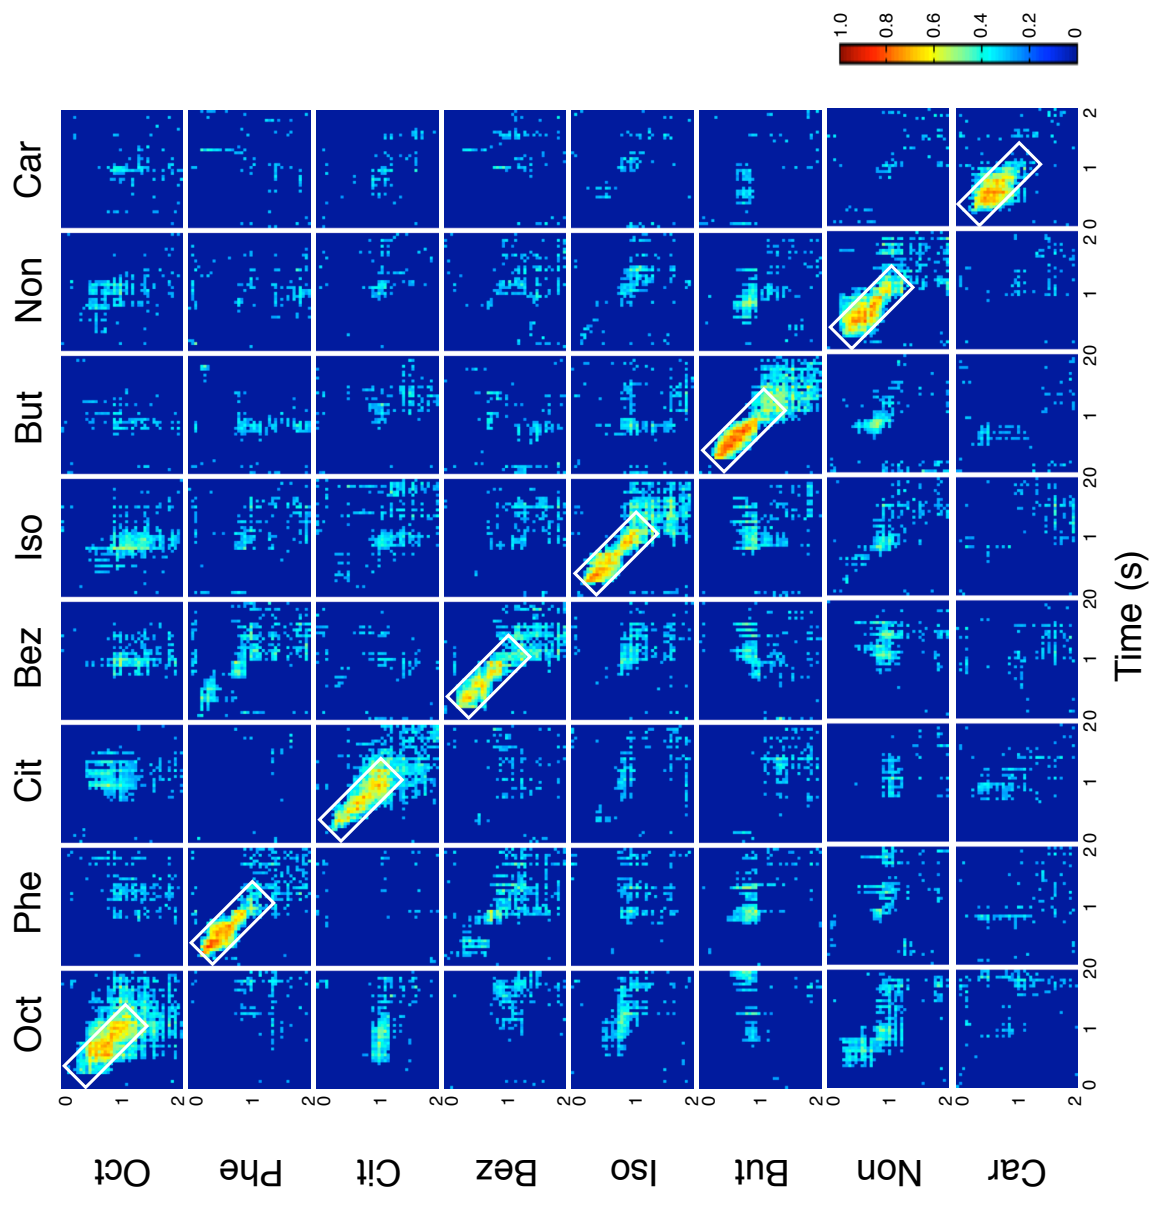


Fig. S2.7




Fig. S2.8


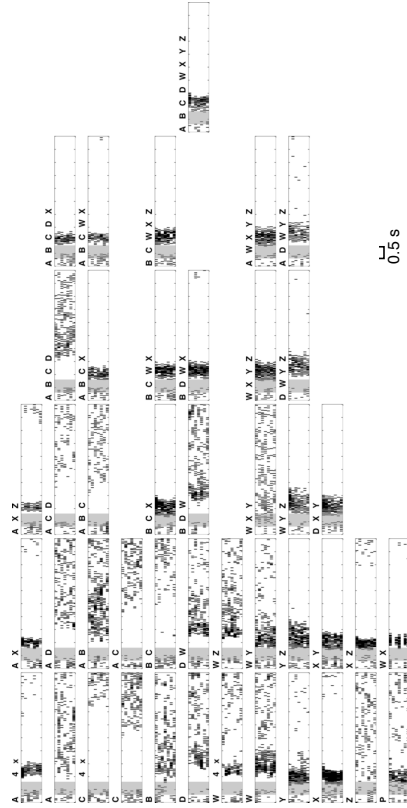
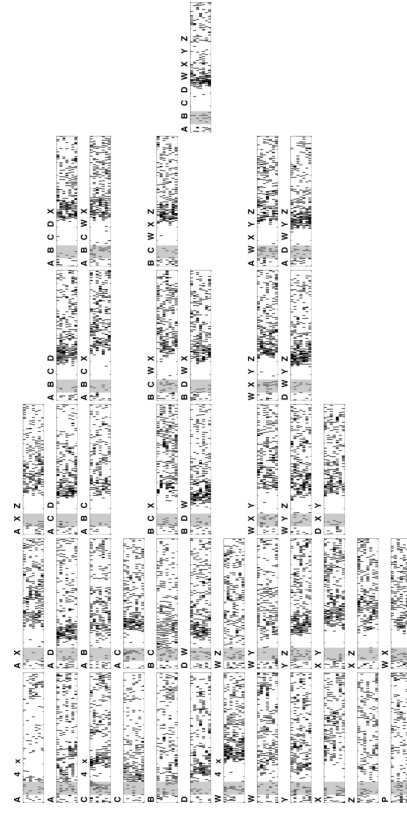


Fig. S2.9

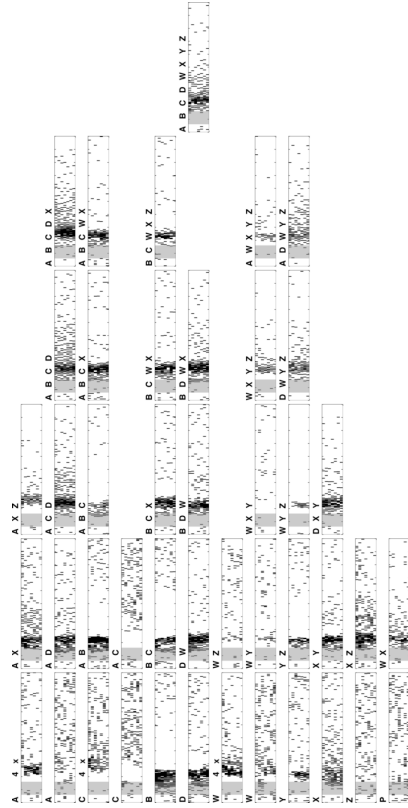
A



B



C



D

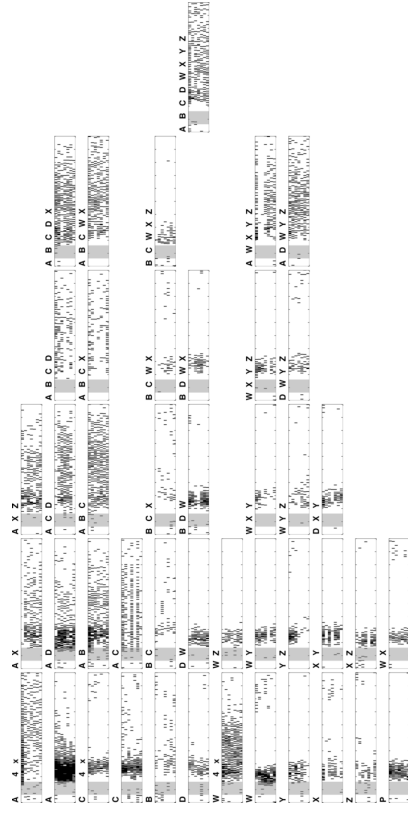


Fig. S2.10


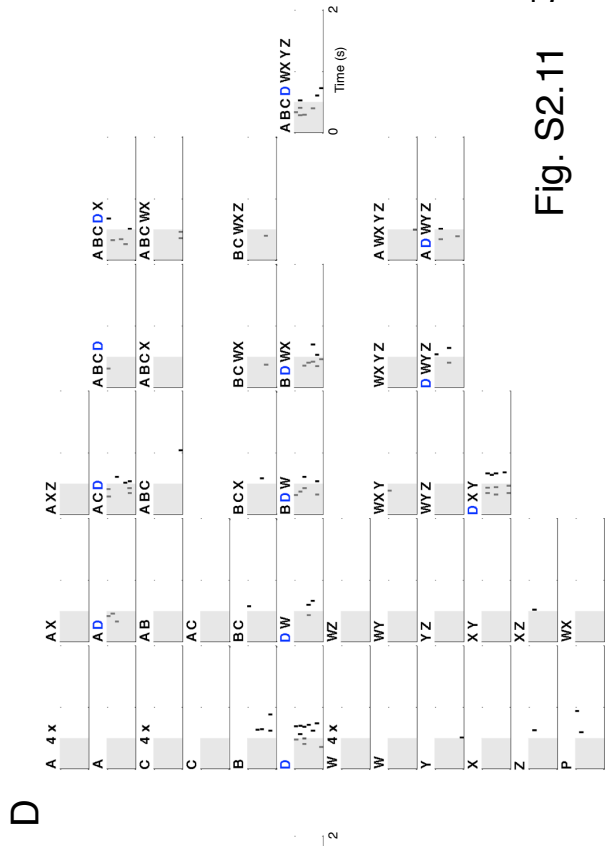

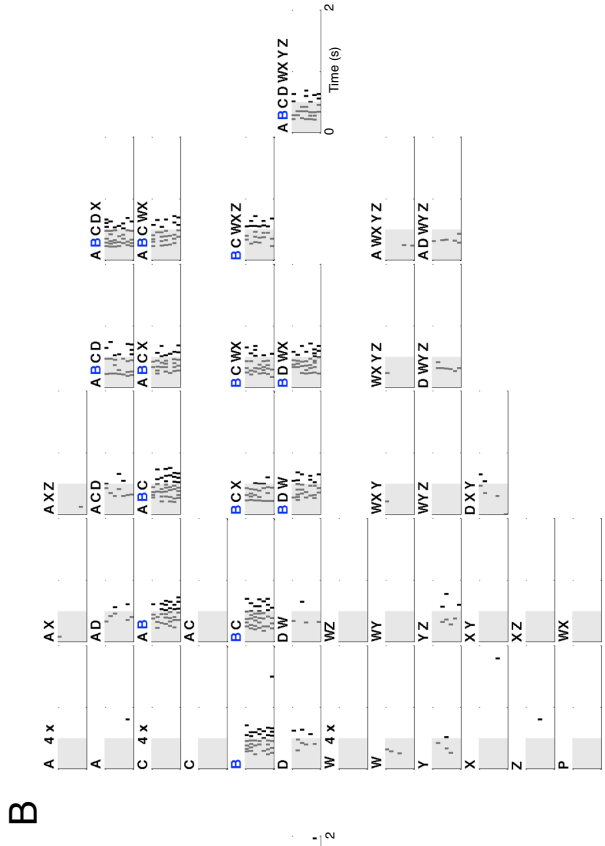


Fig. S2.11 5


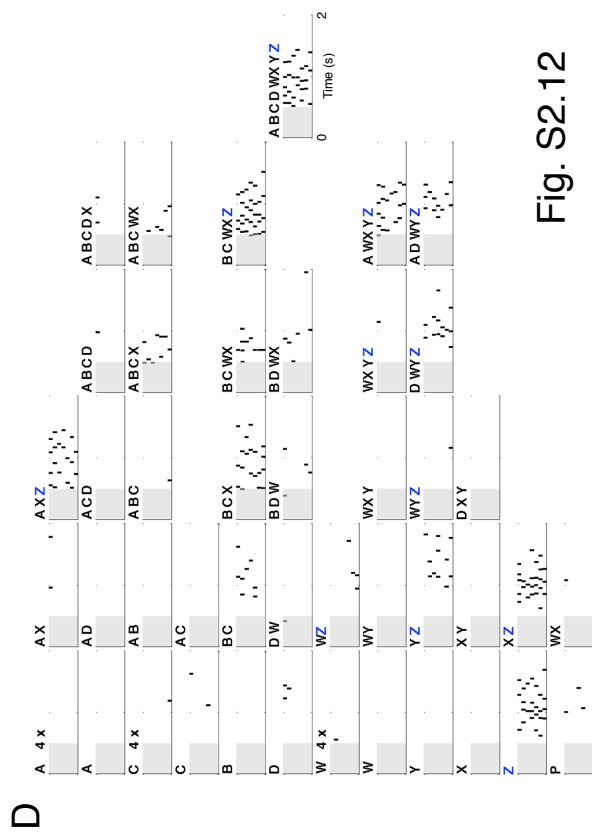
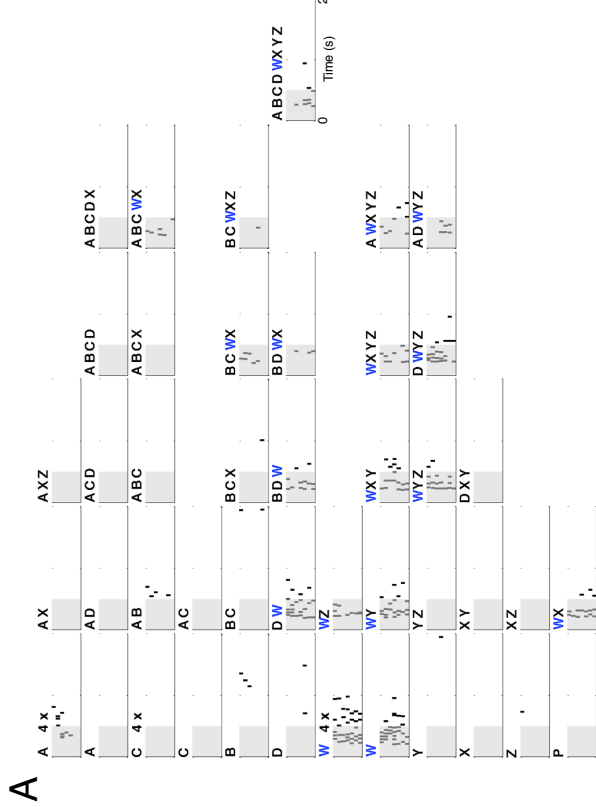
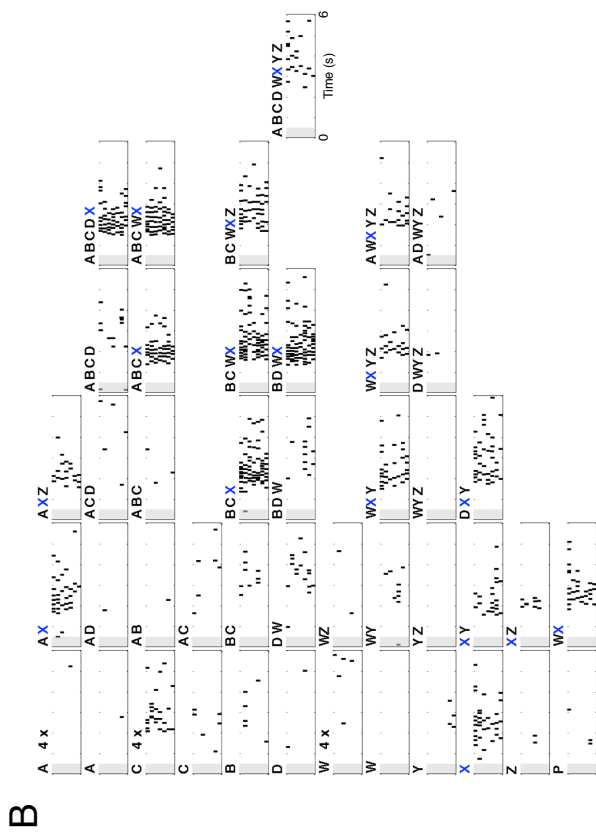


Fig. S2.12

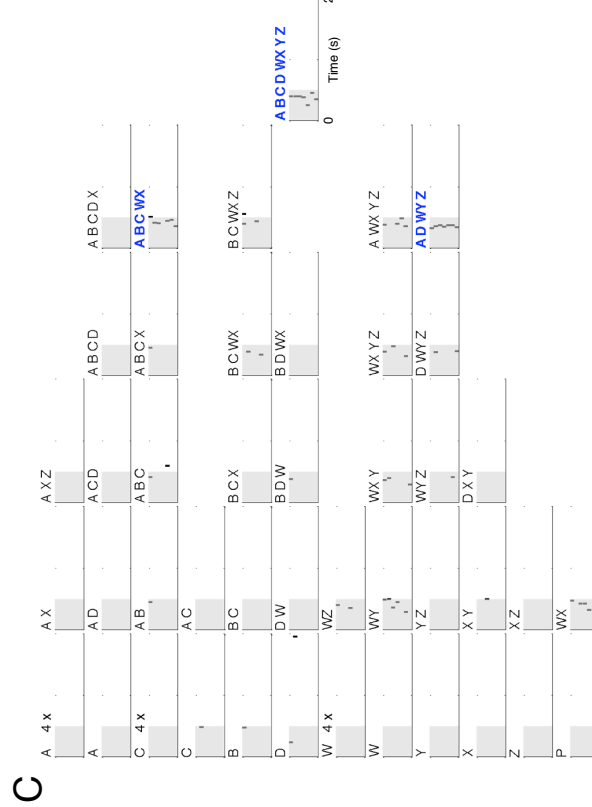
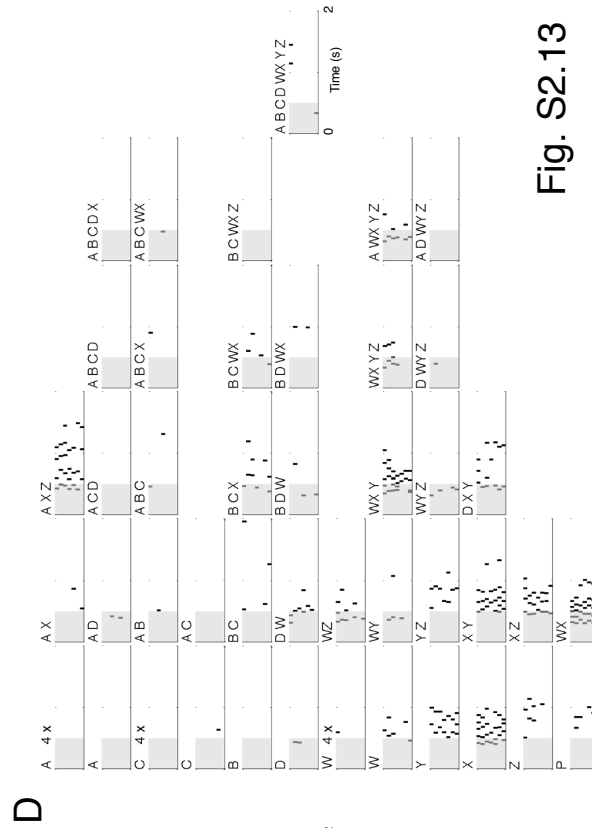

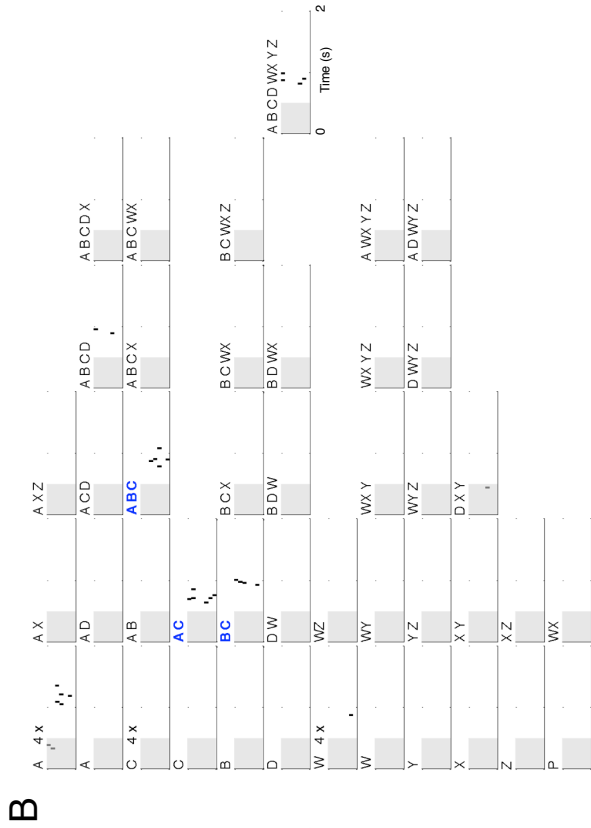


Fig. S2.13 58




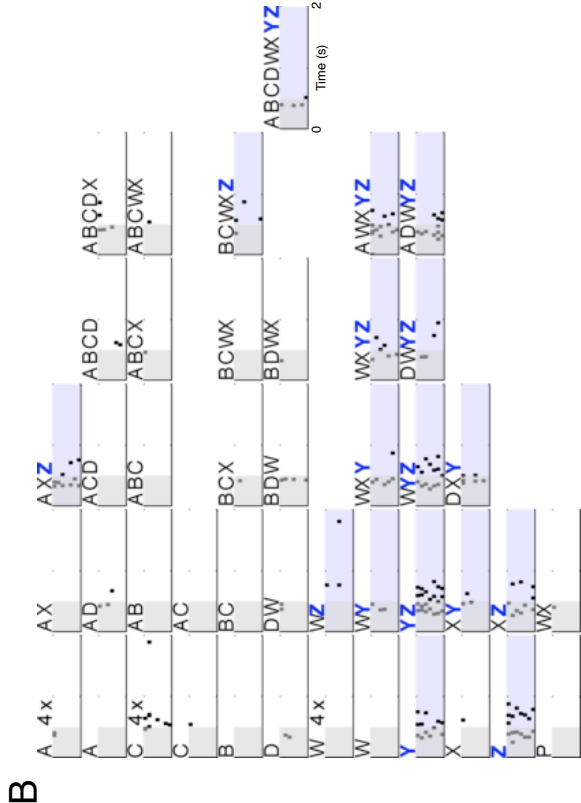



Fig. S2.14

Fig. S2.15






Fig. S2.17





Fig. S2.18

A Binary Mixture Odor Delivery

Fig. S2.19



B Multi-Component Mixture Odor Delivery

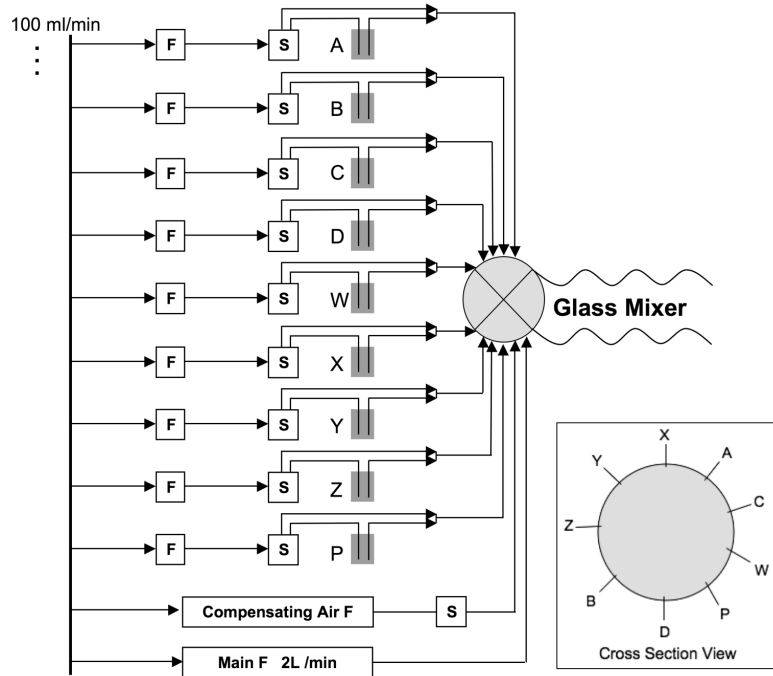




Fig. S2.20

A 175 PNS



B 209 KCs





CHAPTER 2 SUPPLEMENTARY FIGURE LEGENDS

Figure S2.1 Stimulus description for binary mixture experiments.

A–D. Two odors, citral (**A**) and 1-octanol (**B**) were presented individually at six different concentrations, as a morphing mixture from one odor to another in eleven intermediate steps (**C**), and together as a 1:1 mixture at five different concentrations (**D**). Odor pulse was 300 ms, stimulus repeated for 10 trials, each trial was 14s.

Figure S2.2 Stimulus description for multi-component mixture experiments. Eight odor components were presented individually and in combination as 2-, 3-, 4-, 5-, 8- mixtures. Paraffin oil was also presented. In addition, three individual components (1-octanol, citral, isoamyl acetate) were also presented at 4x the component concentration, comparable in concentration to 4- mixtures. Odor pulse was 500 ms, stimulus repeated for 7 trials, each trial was 14s. Right most column (1-, 2-, 3-, 4-, 5-, 8-) indicates n number of components.

Figure S2.3 PN response to binary odor mixtures with ratio- and concentration-specific patterns of excitation and inhibition not predicted by simple arithmetic sum of responses to components.

A–C: Responses of one PN to mixture series that morphes odor citral into odor octanol. Odor pulses are 300 ms; 10 trials per condition.

A. Responses of one PN to six concentrations of pure citral (30, 60, 80, 100, 120, 140 ml/min).

B. Responses of PN to mixture series, from pure octanol (top) to pure citral (bottom). Inserts: overlay of observed (filled PSTHs) and expected (by arithmetic sum of responses to components; black lines) PN responses to morphed mixtures. Mixture conditions are: 140:0, 140:30, 140:60, 140:80, 140:100, 140:120, 140:140, 120:140, 100:140, 80:140, 60:140, 30:140, 0:140 ml/min.

C. Response of PN to six concentrations of octanol. Concentrations as in a.

D–F: Responses of same PN to 1:1 mixtures of the same odors, at different concentrations. **D, F.** same as A, C.

E. Responses to mixture ratios 30:30, 60:60, 80:80, 100:100, 140:140 ml/min.

Figure S2.4 PN trajectories corresponding to odor-morph series.

The same data as in Figure 1C, rotated view. PN population trajectories contain information about both odor identity and the ratio of the binary mixture. The evolution of trajectories with incremental changes in odor ratio is smooth rather than abrupt and there are no apparent discontinuities.

Figure S2.5 Deviation between predicted and experimental PN vectors to odor mixtures

A–B. All distances are computed as correlation distance in 168-dimensional space where each dimension corresponds to the response of 1 PN. Predicted PN population vectors (168–D) are computed from PN vectors in response to odor components

respectively (e.g., 30 Cit:140 Oct is predicted by adding the PN vectors for 30 Cit and 140 Oct; 50 Cit:140 Oct is predicted by adding the PN vectors for 50 Cit and 140 Oct, etc). Deviation between predicted and experimental PN vectors are calculated in each 50 ms bin as correlation distance. Dashed black line show mean instantaneous trial-to-trial distances (distance between one 3-trial average to a different 3-trial average), computed pairwise from the same odor, and averaged across all odor conditions. Shaded region indicates one standard deviation above and below the mean. Within-odor, trial-to-trial distances are high at baseline (since PN trajectories are maximally uncorrelated without stimulus input). With the onset of odor stimulus, trial-to-trial decrease significantly, indicating that trajectories during odor response are less variable than at baseline.

Dashed green line is the mean instantaneous inter-trajectory distance between PN trajectory for Cit (140) and 1:1 mixture (140 Cit:140 Oct), computed pairwise between one 3-trial average and another. Shaded region indicates one standard deviation above and below the mean. Similarly, dashed red line is the inter-trajectory distance between Oct (140) and 1:1 mixture (140 Cit:140 Oct). During the odor response, inter-trajectory distances are greater than mean trial-to-trial distances. Furthermore, the 1:1 mixture trajectory is closer to the trajectory for Oct than for Cit, as indicated by the smaller inter-trajectory distance between Oct and 1:1 Mixture.

How well can we predict PN vectors to mixtures from PN vectors to components? For a short period at odor offset, 0.2–0.3 s in **(A)** and 0.2–0.5 s in **(B)**, the predicted PN vectors do not deviate far from the experimental PN vectors, this is indicated by the fact that these distances are within 1 SD of the mean for trial-to-trial distances. However, for the period of 0.5–0.8 s, the predicted PN vectors deviates from the experimental PN vectors, but these distances are never as large as inter-trajectory distances, suggesting that a linear combination of PN vectors of components leads to a reasonable estimate of true experimental PN vectors for binary mixtures (although they are deviations between the two).

Figure S2.6 Electro-Antennograms (EAG) recordings of single odors.

A–B. EAG recordings from the locust antenna in response to the 8 single odors. Odor concentrations were calibrated to equalize their effectiveness on the receptor array and to ensure that it's operating away from saturation. Methods: an antenna was cut at the

proximal and distal ends, metal electrodes were inserted at both ends; the potential difference was measured by differential amplification. An EAG is a field potential, reflecting the combination of receptor and action potentials (1).

A. Comparison of EAG of single odor concentration (8 single odors: 1-Octanol, Phenetole, Citral, Benzaldehyde, Isoamyl Acetate, 2,3-Butanedione, 2-Nonanone, L-Carvone; 5 trials each; 100 ml/min) in black, with EAG of odors presented at 4 times the single odor concentration (3 single odors: 1-Octanol, Citral and Isoamyl Acetate; 5 trials each; 400 ml/min) in red. Shaded region indicates 1 standard deviation above and below the mean. As shown here, single odors were calibrated to elicit minimal EAG response, at the lower end of its dynamic range. EAG shown for two different antennae (**i and ii**).

B. Comparison of EAG of 8 different single odors, as in (A), and paraffin oil. Concentrations were calibrated to evoke as similar EAG as possible across single odors, so that no one single odor is dominant during odor mixture conditions. Interestingly, paraffin oil evokes no EAG response in isolated antenna, but do can evoke PN and KC responses (see Fig. S10D, Fig. S14C–D). Shaded region indicates 1 SD above and below the mean (5 trials each). EAG shown for two different antennae (**i and ii**).

Figure S2.7 Quantifying differences among PN vectors to different single odors as they evolve over time.

Normalized correlation matrices, between PN vectors in response to single odor conditions. Each of the 8 single odors is compared to every other single odor condition, therefore there are $8 \times 8 = 64$ matrices. The correlation matrices provide a quantitative basis for what is seen qualitatively in LLE (see Figure 1D). For each matrix, each pixel C_{t_1, t_2} represents the color-coded correlation between one 3-trial-averaged 175-PN vector at time t_1 with another at time t_2 , from either the same odor condition (but taken as an average of 3 different trials) or a different odor condition. Thus, a pixel C_{t_1, t_2} represents the correlation of the time slice of trajectory A at time t_1 to the time slice of trajectory B at time t_2 . Diagonal pixels of each matrix represents the correlation between time-matched PN vectors, and each row (of each matrix) represents the correlation from a PN vector at time t_1 to another PN vector at all other times.

Correlations are calculated in 50ms bins. Odor onset is at 0, for 500 ms; PN response onset is ~200 ms later due to delay in odor propagation and olfactory transduction. All

correlations shown are significant ($p < 0.001$). Correlation is initially minimal (dark blue) at baseline until ~ 200 ms after stimulus onset, when PNs begin to respond (yellow, in white box) and lasts for ~ 1.5 s. Blue pixels outside the white box indicate that when PN vectors are not matched in time, correlation is low even when the ensemble response has moved away from baseline, indicating the PN vectors evolve gradually over time. These correlation matrices indicate that each 3-trial average of a single odor condition is highly correlated with another 3-trial average of the same odor condition (as shown by the 8 diagonal matrices), but minimally correlated with different single odors (non-diagonal matrices). Therefore, unique odors are represented by unique PN ensembles over time.

Figure S2.8 Trajectories representing two closely (4 out of 5 part) related mixtures.

Responses of 175 PNs in 50 ms bins (over 4.5 s) to two 4-part overlapping mixtures ADWYZ (red) and AWXYZ (blue) projected in LLE space LLE1-3, consecutive 50 ms bins are linked together to show evolution in time. The two trajectories start to decorrelate at ~ 300 ms after odor onset, and recorrelate after ~ 700 ms (filled circles). Shown are three 2-trial averages for each mixture condition. Open black circle denote baseline. Arrow denote direction of motion.

Figure S2.9 PN response to mixtures not predicted from responses to components

The (smoothed) instantaneous firing rates (black histogram) of one PN (shown as raster in Fig. 2A) in response to all stimulus conditions are compared to those calculated by arithmetic sum (red lines) of their responses to single odor components. To test whether PN responses to mixtures was a simple arithmetic sum of its responses to single odors presented individually, we subtracted mean baseline firing rate and then summed corresponding single PN responses. Firing rates plotted were mean baseline subtracted. Inhibitory responses are therefore represented by the negative of the mean baseline firing rate. Predicted PN responses do not match experimental responses (similar to binary mixtures, see Fig. S3), and deviations between predicted and experimental responses increase with n components of the mixture. Deviations from experimental responses could be in amplitude ('WY') or in time ('DXY') or both ('ABCDWXYZ'). More often than not, we observe that predicted responses were far greater in amplitude than experimental responses, suggesting that a form of gain

control exists in the PN network, whereby combining odors leads to inhibition (see percentage of silent PNs as a function of n , Fig. S19). This form of gain control is most likely mediated by local neurons in the AL network (LN).

Figure S2.10 PN responses to single odors and mixtures.

A–D: Additional PN responses to single odorants and mixtures reveal concentration and mixture-specific interplay of excitation and inhibition. Each row of each panel is a single trial (7 trials from top to bottom). Shown is 5 s. Shaded area correspond to 500 ms odor pulse. Numbers of components organized by column, conditions arranged so that overlapping mixtures are next to each other wherever possible.

Figure S2.11 Component-selective KC responses to single odors and mixtures.

A–D: shows individual KCs selective for unique odor components ('A', 'B', 'C', 'D') and associated mixtures. KCs do not respond perfectly, **(B)** shows a 'B' detector that also responds the mixture 'DWYZ', **(C)** shows a 'C' detector that fails to respond in the 8-mixture condition. Component KC is selective for is indicated in blue.

Figure S2.12 Component-selective KC responses to single odors and mixtures.

A–D: shows individual KCs selective for unique odor components ('W', 'X', 'Y', 'Z') and associated mixtures. **(B)** shows an imperfect 'X' detector that responds 1 s after odor offset, shown is 6 s. **(A, C–D)** shown is 2 s. Component KC is selective for is indicated in blue.

Figure S2.13 Not all KC responses are component-selective, other KC responses.

A. KC that appear to respond to the component 'WX', but not to 'W' or 'X' alone (computation that could be interpreted as 'W' AND 'X').

B–C: Two KCs that respond selectively to specific odor mixtures.

D. KC that responds preferentially to mixtures and sub-mixtures of 'WXYZ', but not to any mixtures or sub-mixtures of 'ABCD'.

Figure S2.14 Not all KC responses are component-selective, other KC responses.

A. KC that responded selectively and only to 4x 'W', (higher concentration of odor 'W').

B. KC that respond to all mixtures containing 'Y' and all mixtures containing 'Z' (computation could be interpreted as 'Y' OR 'Z').

C-D: Two KCs that respond selectively to paraffin oil 'P', but not to any of the odor mixtures, since paraffin oil is odorless and evokes no oscillatory drive in PNs, these could be interpreted as mechanosensory KC responses..

Figure S2.15 KC displayed in the pseudo random order of stimulus presentation.

To ensure units are stationary throughout the entirety of the experiment, certain conditions are presented throughout, and their responses compared qualitatively. Units are discarded if responses to the same odor condition change over the course of the experiment. In the example given, odor conditions 'Y', 'YZ', 'XY', and 'ALL' were presented throughout the experiment (repeat odor conditions indicated by arrows). The order of presentation was left to right, top to bottom. This KC responded very selectively for 'Y'.

Figure S2.16 Comparison of mKC firing rate with and without GGN feedback.

The entire mKC population output is regulated by a negative feedback pathway (GGN). mKC firing rate shown with (right) and without (left) GGN feedback as a function of the number of components n . Dashed lines are experimental KC firing rate, solid lines mKC firing rate. Notice that without GGN feedback, mKC firing rate (left) increases at much faster rate (as a function of n) than if GGN feedback is implemented (right). With GGN activation at ~400 units (see Methods), KC firing matches mKC firing.

Figure S2.17 Model KC population firing rate matches experimentally observed KC firing rate.

The red curve is the mean firing rate of a population of 10,000 mKCs in the one-second following odor onset, as a function of the number of odor components present in the odor. The dashed gray curve is the experimentally observed curve.

Figure S2.18 PN and KC population statistics.

A–C: All statistics were computed from 175 PNs and 209 KCs; each presented seven times with 44 different stimuli. Silent cells were computed with consecutive, non-overlapping 100 ms bins; responsive cells were computed with consecutive non-overlapping 50 ms bins; firing rate was computed by convolving 10 ms binned spikes with a 20 ms width Gaussian filter. Thin black lines show averages for each odor condition, thick lines are averages across odor conditions. Blue bar is 500 ms odor pulse, 3 s shown. Horizontal line in (A–B). indicate maximum or minimum value reached during the single component conditions.

A. Mean PN and KC firing rate as a function of the n number of odor components in the mixture. Mean firing rate for PNs remains approximately constant with increasing n (with a slight decrease). In comparison, KC mean firing rate clearly increases as a function of n . Interestingly, PN firing for single components at 4x concentration is higher than 1x, but no significant differences were detected in KC firing rate.

B. Percentage of silent cells for each time bin. A cell is defined as silent during a 100 ms time bin if it fired no spikes in that time bin in all 7 stimulus trials. Notice that percentage of silent PNs increases as a function of n , indicating increased inhibition from LNs, as a form of gain control on the output of the PNs. The peak of the percentage of silent PNs is reached ~200–300 ms later than the peak of PN firing (see C). Four part odor mixtures elicited greater inhibition than single components at comparable concentration, suggesting that the form of gain control observed here is specialized for mixtures. In comparison, at baseline, all KCs are silent, this dips to ~90% for single components and is quickly shut completely within 500 ms of odor onset, from the PN–KC model, we attribute to feedback gain mechanism of the GGN which keeps KCs sparse.

C. Percentage of responsive PN- and KC- odor pairs as a function of n odor components. Very few PNs and KCs were responsive at baseline by our measure. However, shortly after odor onset, 7–14% of PNs are responsive, in comparison to 0.5–1% of KCs in any 50 ms time bin. The number to the upper right of each panel shows the cumulative proportion of responding cells over 3 s.

Figure S2.19 KC population activity matches well with PN synchronization as revealed by the local field potential (LFP).

A. Larger odor mixtures (as n components in mixture increases) elicit greater oscillatory power in the local field potential (LFP). Single odor components (at 1x concentration) elicit minimal (if any) power in LFP (left most column). Shown single trial (from 1 experiment) to all odor conditions. LFP was band-pass filtered, 5–35 Hz.

B. Mean normalized LFP power in 5–35 Hz band (mean over all odors and 6 locusts, 200 ms sliding window in 50 ms steps) in black line (SD in gray). Mean instantaneous KC firing rate (10 ms binned spikes convolved with 20 ms Gaussian) is superimposed in red. Increases in both KC firing rate and LFP power are well matched. Notice that increases in LFP power with n components in the mixture cannot be explained by concentration alone, as 4-part odor mixtures (4-) elicit still greater power than single odors at 4x the concentration (4x 1-) (the two are at equivalent concentration and elicit equivalent EAG responses, not shown). Interestingly, Paraffin oil does not elicit power in the LFP or EAG responses (see Fig. S6), but does lead to increases in KC firing rate. Odor bar was shifted by 100 ms in B to better align with true odor onset; Odor bar not shifted in A.

Figure. S2.20. Custom designed olfactometer for mixture experiments

A–B: An olfactometer was constructed using small diameter (1/32" inner diameter) Teflon tubing and compression fittings to minimize dead space and delay times. Flow rates of independent air streams were independently controlled by mass flow controllers (range 20–200 ml/min). By mixing odorized air streams with defined flow rates, different mixture ratios were achieved.

A: Binary mixture experiments; B: Multi-component mixture experiments.

Figure. S2.21. Responsive PNs and KCs.

A–B: Excitatory responses (filled squares) of individual PNs and KCs (columns) (n=175PNs, 209KCs) to 44 different odor stimuli. Open squares denote inhibitions (PNs) or absence of a response. Responsiveness was defined as responding with at least 1 spike in each trial in 4/7 trials and above 1.5 SDs of baseline (see Methods).

Figure. S2.22. Number of ‘exclusive’ KCs that responded to single odors and mixtures.

Responses properties for “exclusive” KCs (see Methods), those that responded to only 1 of 8 odor components, and whether they responded to in-class and out-class mixtures.

References:

1. B. M. Broome, V. Jayaraman, G. Laurent, *Neuron* **51**, 467 (Aug 17, 2006).

2.3 Methods

2.3.1 *Preparation and stimuli*

Results were obtained from 61 locusts (*Schistocerca americana*) in a crowded, established colony. We recorded from 168 PNs for the binary mixture experiments, 175 PNs for the multi-component mixture experiments and 209 KCs from 13 (42 groups), 11 (36 groups) and 37 locusts (53 groups) respectively. Experiments were typically conducted using left and right ALs and MBs in each animal. Young adults of either sex were immobilized, with one or two antennae intact for olfactory

stimulation. The brain was exposed, desheathed and superfused with locust saline, as previously described (42). Odors were delivered by injection of a controlled volume of odorized air within a constant stream of dessicated air. Teflon tubing was used at and downstream from the mixing point to prevent odor lingering and cross-contamination. Due to different odor stimuli requirements between the binary morphing and the multi-component mixture experiments, two independent odor delivery systems were built.

2.3.2 Binary-mixture experiments

Two odorants (1-octanol and citral, Sigma) were stored as pure solutions in independent 500ml custom-made bubblers. The odor nozzle (1 cm diameter, Teflon) was placed 1 cm from the antenna and supplied a constant 1 l/min carrier stream of desiccated, filtered air. The flow of each odor was controlled by an independent electronic flow meter with feedback control (McMillan Model 80-D) upstream of the bubbler and a solenoid placed downstream. Relative odor concentrations were varied by controlling flow through each bubbler (30, 50, 80, 100, 120 and 140 ml/min). These concentrations spanned the dynamic range of electro-antennal responses recorded in recordings from isolated antennae (electro-antennograms, or EAGs). A large vacuum hose placed behind the antenna guaranteed the quick removal of odorants from the space

surrounding the antenna. Odor puffs were triggered automatically using a custom computer interface (LabView, National Instruments Inc.). Trials were 14 s long, with 300-ms-long odor puffs presented 2 s after trial onset; each odor condition was repeated 10 times. Stimulus description shown in Fig. S2.1.

2.3.3 Multi-component-mixture experiments

Individual odors were chosen to be chemically different and their respective concentrations adjusted to ensure that no one odor dominated over the others due to intrinsic differences in vapor pressure. In practice, odor concentrations were calibrated by dilution in paraffin oil to equalize EAG responses, recorded from isolated antennae. Paraffin oil alone elicited negligible EAG response. The odors were: 1-octanol (A), diluted 0.7 ml/10 ml; phenetole (B), diluted 0.15 ml/15 ml; citral (C), pure 10 ml; benzaldehyde (D), diluted 0.02 ml/15 ml; iso-amyl-acetate (W), diluted 0.1 ml/10 ml; 2,3- butanedione (X), 0.04 ml/15 ml; 2-nonanone (Y) diluted 2 ml/15 ml; L-carvone (Z), pure 10 ml; Paraffin oil (P), pure 15 ml. The individual odors were each placed into a glass vial (60 ml). The headspace content was carried by puffs of desiccated and filtered air, with a flow rate of 100 ml/min for individual odors and 400 ml/min for paraffin oil. Three odors: 1-octanol (A), citral (C), and iso-amyl-acetate (W) were also presented at a second, higher concentration, by increasing

flow rate to 400 ml/min. Odor mixtures were presented by combining the single odorants. For example, AB was the combination of 100 ml/min of A with 100 ml/min of B, with a total odor flow of 200 ml/min; thus, total odor concentration was higher during mixture conditions. A compensating stream of desiccated air was used to ensure that total air flow remained constant throughout the experiment. The odors were mixed in a custom-built corrugated glass tube (~15 cm long, 1 cm diameter), with a total flow of 2 l/min to ensure turbulent mixing. The individual odor lines were arranged along the circumference of the mixer. Trials were 14 s long, with odor puffs presented for 500 ms, 2 s after trial onset, and repeated 7 times with each stimulus. To minimize the potential effects of priming (Backer, 2002), single odorants were presented first, followed by 2-, 3-, 4-, 5- and 8- mixtures. The order of presentation within each mixture group was pseudo-random. Stimulus description shown in Fig. S2.2.

2.3.4 Electrophysiology

Two types of tetrodes were used for extracellular recordings. Silicon probes were obtained from NeuroNexus. Wire tetrodes were constructed with insulated 0.0005" and 0.0004" wire (REDIOHM wire with PAC insulation). Four strands of wire were twisted together and heated to partially melt the insulation. The tip was cut with fine scissors and each

channel tip was electroplated with a gold solution to reduce the impedance to between 200 and 250 k Ω at 1 kHz. The same custom-built 16-channel preamplifier and amplifier were used for both types of tetrodes. Two to four tetrodes were used simultaneously. The preamplifier had a gain of 1, and the amplifier gain was set to 10,000x. Because of low baseline activity and low response probability in KCs, fewer KCs than PNs were usually isolated in a typical recording session. Tetrodes were placed within the AL or MB soma clusters, peripheral to the neuropils at depths between 50 and 200 μm . For some MB recordings (KCs, LFP), probes were pressed on the surface of the MB. Cell identification was unambiguous because PNs are the only spiking neurons in the locust AL—LNs do not produce sodium action potentials (72)—and because all the somata located dorsal to the MB calyx belong to KCs. Recording locations were tested randomly across the MB and selected if activity could be elicited by any of the 44 odor conditions. Identical stimuli were presented at the beginning, middle and end of the experiment to check that clusters had not drifted significantly over the course of the experiment. Drift was estimated qualitatively by determining if a given neuron's responses to each odor were similar across these three sampling periods. Hints of drift then led to examination of the waveform clusters.

2.3.5 *Recording constraints and potential sampling*

biases.

Because PNs respond very promiscuously to odors (53), no effort was made to find PNs that responded to our stimuli. As soon as good signals suggestive of separable PN clusters could be seen, recordings started and responsive PNs were always found. Our estimates of PN response-probabilities are therefore likely close to true values.

By contrast, KCs respond very sparsely to odors and their individual baseline activity is ~1 spike every 30 s on average (53). Hence, significant effort was made to find KCs that responded to some at least of the odors in our panel prior to initializing an experiment. Due to the large number of conditions in our experiments, we could not pre-test all 44 stimuli. Rather, we searched for responsive KCs by presenting the eight monomolecular odors; we selected a recording position from which some spikes could be recorded in response to any one of these stimuli. Due to the rarity of KC spikes, KC-spike cluster models were defined using all trials (usually ~50 conditions, 7 trials each, 14 s per trial). The condition in the middle of the set was used to calculate the noise covariance matrix (84). The threshold was set typically at 4–5 times each channel's signal SD. The model generated by this method was refined using criteria identical to those used with the PN data. Stability over the course of the experiment was assessed after sorting and was based on a stable

baseline firing rate over the course of the experiment. Whereas experiments started when some signal deemed hopeful had been detected on at least a couple of channels of one tetrode, not all the KCs that we characterize here were recorded because they had been identified during the initial selection process: many happened to be recorded fortuitously, and isolated *post hoc*. Yet, because our recordings started with an active search for responding KCs, they are probably biased towards cells selective for one of our primary odors. This observation does not change anything to our conclusions, except in a quantitative sense: the true response-probabilities of KCs must be even lower than we presently estimate (Fig. S2.17).

2.3.6 *Extracellular data analysis*

Tetrode recordings were analyzed as described in (84). Briefly, data from each tetrode were acquired continuously from the four channels (15 kHz/channel, 12 bit/sample), filtered (custom-built amplifiers, band-pass 0.3–6 kHz) and stored. Events were detected on all channels as voltage peaks above a pre-set threshold (usually 2.5–3.5 times each channel's signal SD for PNs and 4–5 SDs for KCs). For any detected event on any channel, the same 3-ms window (each containing 45 samples) centered on that peak was extracted from each one of the four channels in a tetrode. Each event was then represented as a 180-dimensional

vector (4×45 samples). Noise properties for the recording were estimated from all the recording segments between detected events, by computing the auto- and cross-correlations of all four channels. A noise covariance matrix was computed and used for noise whitening. Events were then clustered using a modification of the expectation maximization algorithm. Because of noise whitening, clusters consisting of, and only of, all the spikes from a single source should form a Gaussian ($SD = 1$) distribution in 180-dimensional space. This property enabled us to perform several statistical tests to select only units that met rigorous quantitative criteria of isolation (84).

2.3.7 *Analysis*

Projection neurons

MATLAB and the Statistics Toolbox (The MathWorks, Inc.) were used for data analysis and modeling. To test whether a PN response to an odor mixture was a simple summation of its responses to the odors presented individually, we convolved all spikes in the responses with a Gaussian of width 20 ms, averaged the smoothed spike counts across trials, and subtracted mean baseline spike counts (calculated from data preceding stimulus onset). We then compared the response to the mixture

condition with the arithmetic sum of the responses to odors presented individually.

For nonlinear dimensionality reduction with Locally Linear Embedding (85), we used code from Sam Roweis (<http://www.cs.toronto.edu/~roweis/lle/>) with Gerard Sleijpen's code for the JDQR eigensolver (<http://www.math.uu.nl/people/vorst/JDQR.html>). In the figures shown for nonlinear dimensionality reduction with LLE, we used as inputs 168-D time slices, 50 ms long each, averaged over 3 trials (binary mixtures), and 175-D time slices, 50 ms long each, averaged over 2 trials (multi-component mixtures). Other details are as described in (7).

PN and KC probability of response

We analyzed spiking data for a 1 s window shortly after odor onset, in all multi-component odor mixture conditions. A PN or KC was classified as responding if its firing behavior during the 1-s window met two independent criteria of response *amplitude* and *reliability*.

Amplitude: the neuron's firing rate (measured in successive 200-ms bins, averaged across all trials) had to exceed the mean baseline and n standard deviations of the baseline rate (measured across 200ms bins and all 7 trials) in at least one bin within the response window. Baseline

rate was measured for each cell-odor pair over a period of 600 ms preceding stimulus onset and over all 7 trials. Values of n of 1.5 or 2 gave low rates of both false positives (during baseline) and false negatives (during stimulation) in PNs. Values of n between 0 and 4 made no significant difference with KCs. We show results with $n = 1.5$.

Reliability: to ensure that responses detected were reliable even at low firing rates in KCs, we required that at least one trial more than 50% of all trials (i.e., at least 4/7) with each odor contained at least one spike during the response window. Our metric for responsiveness is extremely conservative, because it measures PN activity for only 1s shortly after odor onset. In reality the dynamics of PNs last for as long as 3-4s after odor offset (e.g., rebound excitation that occurs later in PNs is not captured by our metric). Responsive PNs and KCs by this metric are shown in Fig. S2.17.

To assess whether PN and KC responses generalized to odor mixtures, we calculated the conditional probability of response $Prob_C$. Given that a cell responded to an odor, we then calculated $Prob_C$ to mixtures containing that odor component (in-class), as a function of the n number of components in the mixture, (1-, 2-, 3-, 4-, 5- and 8- mixtures). First, we found responsive PNs and KCs (defined by the metric above); if a cell responded to an odor component (e.g., odor A) the $Prob_C$ to mixtures was defined as:

$Prob(n\text{- mixtures containing A} \mid \text{odor A}) = (\text{number of } n\text{- mixtures containing odor A that cell responds to}) / (\text{number of } n\text{- mixtures containing odor A})$

Similarly, we also calculated $Prob_c$ to mixtures not containing the odor component (out-class):

$Prob(n\text{- mixtures not containing A} \mid \text{odor A}) = (\text{number of } n\text{- mixtures not containing odor A that cell responds to}) / (\text{total number of } n\text{- mixtures not containing odor A})$

This procedure was repeated for all cell-odor pairs across all n - mixtures (1-, 2-, 3-, 4-, 5-, and 8-), and we averaged across all cell-odor pairs as a function of the number of odor components in the mixture. Shown in Fig. 2.2d are $Prob_c$ plotted as a function of the number of components in the mixture; the shaded region represent the 30-70% percentile distribution of all cell-odor pairs. Fig. S2.21-2 shows the breakdown for different odor components. PN $Prob_c$ for in-class mixtures was 1, 0.65, 0.62, 0.58, 0.57 and 0.57 for 1-, 2-, 3-, 4-, 5-, and 8- mixtures respectively. PN $Prob_c$ for out-class odors and mixtures were 0.34, 0.44, 0.43, 0.45, 0.44 for 1-, 2-, 3-, 4-, and 5- mixtures respectively. There were no “out-class” 8- mixtures since the 8- mixture contains every odor component. By comparison, KC $Prob_c$ for in-class mixtures was 1, 0.60, 0.63, 0.54, 0.52, and 0.54 for 1-, 2-, 3-, 4-, 5-, and 8- mixtures

respectively. KC $Prob_C$ for out-class odors and mixtures was 0.16, 0.23, 0.27, 0.27, and 0.26 for 1-, 2-, 3-, 4-, and 5- mixtures. Thus, although $Prob_C$ to in-class mixtures were similar for PNs and KCs, $Prob_C$ to out-class mixtures were much lower for KCs than PNs, indicating a greater separation of in-class to out-class responses for KCs. In addition, for KCs, we also calculated the probability of response $Prob_{C,unique}$ to mixtures conditioned on unique responses to components, an “exclusive” computation. That is, we first found KCs that responded to only one particular odor component and calculated their response probabilities to mixtures (as before). KC $Prob_{C,unique}$ for in-class mixtures were 0.44, 0.45, 0.38, 0.38, 0.3 for 2-, 3-, 4-, 5-, and 8- mixtures respectively. KC $Prob_{C,unique}$ for out-class mixtures were almost negligible at 0, 0.04, 0.08, 0.09 and 0.07 for 1-, 2-, 3-, 4-, and 5- mixtures respectively. This shows that out dataset contained a group of KCs with very selective responses to single odors and in-class mixtures. We could not apply this metric to PNs for there were no PNsthat responded only to 1 component. We analyzed this further by counting the number of KCs that responded to single odor components and mixtures (Fig. S2.23). Out of 209 KCs, 47 KCs responded uniquely to one odor component, more than half of these KCs, 27 responded to at least one of 4- and 5- in-class mixtures, and 14 responded to 8- mixture of all single odor components. About a quarter of these KCs responded to all in-class mixtures across n number of components.

ROC analysis

With each responsive neuron, we measured the performance of an ideal observer (using firing rate across trials) at predicting whether or not an odor component was present in the stimulus. ROC analysis was developed in signal-detection theory to express the tradeoff between the hit-rates and false-alarm-rates of classifiers (81). Given a classifier and an input class, there are four possible outcomes. If the input class is positive and it is classified as positive, it is counted as a true positive. If the input class is negative and it is classified as positive, it is counted as a false positive. ROC curves depict relative tradeoffs between true positives (TP rate) and false positives (FP rate).

		True Class of Odor Stimuli		
		P	N	
Classification by PN or KC	P	True Positives	False Positives	Precision = $\frac{TP}{TP+FP}$
	N	False Negatives	True Negatives	Accuracy = $\frac{TP+TN}{TP+FP+TN+FN}$
		$TP\ rate = \frac{TP}{TP+FN}$	$FP\ rate = \frac{FP}{FP+TN}$	

How reliably can the activity of a PN or KC tell us whether or not an odor component occurred in an odor mixture? We first found PNs and KCs that respond (see response metric above) to odor components. We then

evaluated each responsive neuron as a classifier for a particular odor component (e.g., A). We partitioned all odor conditions into two classes: the in-class of all odor conditions including A (A-high, A, AB, AC, AD, AX, ABC, ACD, AXZ, ABCD, ABCX, ABCDX, ABCWX, ADWYZ, AWXYZ, ALL) and the out-class of all conditions without A (Chigh, Whigh, B, C, D, W, X, Y, Z, BC, DW, WX, WY, WZ, XY, XZ, YZ, BCX, BDW, DXY, WXY, WYZ, BCWX, BDWX, DWYZ, WXYZ, BCWXZ). This was repeated for all responsive cells-odor class pairs. We summed spikes across all trials within a 1s window shortly after odor onset, and used a sliding threshold to determine the probability of true positives (among all in-class conditions) and the probability of false positives (among all out-class conditions). Decreasing the threshold increases the probability of true positives but also of false positives. The ROC curves for each cell-odor class pair are shown in Figure 2.2e. A cell responding indiscriminately to odors and mixtures would have an ROC curve close to the diagonal; a cell responsive to all odor mixtures containing one particular component but not to mixtures containing other components would have a convex ROC curve far from the diagonal. The area under the curves is therefore a measure of the ability of each neuron to classify successfully the presence of an odor component. A value of 0.5 corresponds to chance performance; a value of 1 indicates perfect accuracy in classification. The distribution of the areas under the curves across all cell-odor class pairs was centered at 0.57 ± 0.18 for PNs (significantly shifted to the right of

0.5, $p < 6 \times 10^{-10}$, Wilcoxon signed rank test) and 0.73 ± 0.16 for KCs (significantly shifted to the right of 0.5, $p < 4 \times 10^{-17}$, Wilcoxon signed rank test). The distribution for KCs were significantly shifted to the right of that for PNs ($p < 3 \times 10^{-13}$, Wilcoxon rank sum test). For PNs and KCs that responded to multiple single odors, we plotted the ROC curve of each of the primary odor classes that they responded to. Given the nature of the partitioning of in- and out- classes, if a PN or KC was a good classifier for one odor component, then it was by necessity not a good classifier for another odor component. We repeated the ROC analysis for longer odor windows of 1.4 s and 2 s, and for single trials and observed no differences.

PN-KC model

KC model summary

Kenyon cells were modeled as leaky integrate-and-fire units receiving feed forward excitatory input from selected PNs and feed forward inhibitory inputs from all LHIs. These potentials were summed to give the membrane potential of the model KC (mKC). If this potential exceeded a threshold value θ , a spike was elicited and the cell entered an absolute refractory period, representing the after-hyperpolarization (AHP) of KCs (42). Feedback inhibition from the giant GABA-ergic neuron (GGN,

Papadopoulou et al). was modeled as an increase in the firing threshold of all KCs, based on the instantaneous KC population activity. Details of the model are given below.

PN Input to the model

The experimentally recorded spiking activity of 151 PNs with simultaneously recorded LFP in response to seven trials of 44 different odor conditions served as the input to the model. Each model trial was 5 seconds long, consisting of 1 s of pre-odor baseline activity, 500 ms of odor presentation and 3.5 s of post-odor response. The set of 151 PNs was formed by pooling PN data over several experiments. To approximate simultaneous recording conditions, all spikes were aligned to the same LFP by linear phase-warping.

Phase-warping of PN spike output

To generate large PN population vector appropriate for driving simultaneously our mKCs, we aligned the spikes of PNs recorded in different experiments to the same clock; we used the LFP recorded with each experiment as the common time reference (53). Each PN spike time was thus assigned a phase relative to its own LFP oscillation, and all

spikes and PN outputs were then assigned a time and phase, based on a common 20Hz clock. Practically, raw LFP signal from a mushroom body electrode was band-pass filtered from 10–30 Hz (non phase-distorting Butterworth filter, built-in Matlab functions). To compute spike phase histograms, each PN spike was assigned a phase with respect to the simultaneously recorded LFP using a simple algorithm. First, all peaks and troughs that exceeded 0.5 SD of baseline (non-odor evoked) fluctuations were detected in the band-pass filtered LFP. Spikes were assigned a phase if they fell between a peak and a trough (in either order), and if less than 35 ms separated the peak and trough. Phase was then assigned based on a linear scaling of phase values between a peak and trough (0° – 180°), or trough and peak (180° – 360°). Spikes with no phase attribution were distributed uniformly across all phases.

Direct excitatory inputs

The excitatory inputs to mKCs start with the spiking activity of 55 PNs, selected randomly out of the population of 151 recorded PNs. Spikes were convolved with a first order EPSP filter characterized by a gain parameter A_{pn} and a decay parameter τ_{pn} . A_{pn} and τ_{pn} were updated according to the instantaneous membrane voltage V_m of the cell,

expressed as a fraction of the baseline threshold θ_0 of the mKC. EPSP sharpening (53) was implemented according to the rules below:

$$A_{pn}(t) = \begin{cases} 1 & \text{if } V_m(t-1)/\theta_0 < 0.5 \\ 4V_m(t-1)/\theta_0 - 1 & \text{otherwise} \end{cases}$$

$$\tau_{pn}(t) = \begin{cases} 15 & \text{if } V_m(t-1)/\theta_0 < 0.5 \\ 25V_m(t-1)/\theta_0 - 20 & \text{otherwise} \end{cases}$$

Note that the value of τ_{pn} was only changed if a new value was lower than the old value.

Feed-forward inhibitory inputs

Feed-forward inhibition of KCs is via a small population of lateral horn interneurons (LHIs), that show poor odor tuning (53). Hence, the LHI population was modeled as a single source of inhibition, driven by the spiking output of all 151 PNs and delayed by Δ_{lhi} ms relative to PNs to model the PN- LHI phase lag of just under one half of one oscillation cycle (53). This output was convolved with a constant first-order IPSP filter with a gain of A_{lhi} and a decay of τ_{lhi} . The resulting IPSPs were sent to all KCs simultaneously.

mKC membrane voltage, threshold, and spike generation

E- and I-PSPs were summed at each mKC to give the instantaneous membrane voltage V_m :

$$V_m(t) = \begin{cases} EPSP(t) + IPSP(t) & \text{if } EPSP(t) + IPSP(t) > 0. \\ 0 & \text{otherwise} \end{cases}$$

If V_m exceeded the instantaneous firing threshold θ , a spike was elicited and the mKC entered an absolute refractory period of duration τ_{ref} [Laurent and Naraghi 1994].

Threshold and feedback inhibition

The threshold of the cell is defined as

$$\theta(t) = \theta_0 + \theta_{ggn}(t)$$

θ_{ggn} reflects all-to-all feedback inhibition provided by the GGN and is defined as [Papadopoulou et al.]

$$\theta_{ggn}(t) = K_{ggn} \left(\frac{1}{N} \sum_{i=1}^N S_i(t-1) \right) * Filt_{ggn}$$

Here N is the total number of KCs in the population, S_i is a binary value indicating whether the i^{th} KC produced a spike, K_{ggn} is a constant gain

parameter, Filt_{ggN} is an FIR filter modeled after the GGN impulse response derived from experiment, and the star represents linear convolution.

Connectivity and bias procedure

Each run of the model consisted of 151 PNs connected to a population of 10,000 Kenyon cells so that the KC to PN ratio matched the anatomically observed one (59). In the 'no-bias' condition, each KC was connected to a randomly selected subset of 55 PNs among the 151 available. This fan-in fraction to each KC was within the experimentally observed range of 35 – 65 % (59). Baseline threshold θ_0 and GGN gain K_{ggN} were tuned so that the mean KC population firing rate as a function of the number of odor components present in a mixture matched the values recorded in our experiments experimentally. K_{ggN} is one of the most critical components of the model to ensure a compressed range of KC response probabilities. One hundred such mKC populations were run, for a total of 1,000,000 model KCs. Conditional response probabilities were computed for each population, and were observed to be a poor match to the experimental distributions.

Connectivity bias was employed to improve the segmentation performance of the KC population. The 1,000,000 model KCs were

filtered for KCs whose individual conditional probability distributions matched those of the biological observed population, namely a minimum probability of in-class response of 0.3, and a maximum out-class response probability of 0.2. About 1,200 such cells were found in the set of 1,000,000 model KCs. This population of 1,200 mKCs was partitioned into 8 groups according to which single odors they responded to. Note that these groups were not disjoint: an mKC could respond to multiple single odors. For each of the 8 groups, the connectivity vectors of the member mKCs were overlaid to produce a frequency distribution of connectivity to each PN. These distributions were then sorted. The first PN in the sorted vector for odor A was the PN most frequently connected to by mKCs segmenting A, and so on.

To implement the bias, a new population of KCs was generated. The connectivity of a fraction bf of this population was biased, while the remaining mKCs were assigned random connectivity as before. The biased fraction of the population was partitioned into 8 disjoint groups, one group for each odor component. The first mh of the 55 inputs to the KCs in each group were set to be the first mh PNs in the PN popularity vector for the associated odor. For example, an mh level of 0 corresponded to no bias, an mh level of 1 meant that the connectivity vector for a biased mKC contained the most popular PN for the odor the mKC is meant to segment, etc. The remaining $55 - mh$ inputs to a mKC

were set randomly. The values of bf and mh were swept from 0 to 1, and from 0 to 15, respectively, and it was found that $bf = 0.2$ and $mh = 15$ produced the best match to the experimentally observed conditional response probabilities.

Summary of model parameters

Parameter	Description	Value
N	Number of KCs per model run	10000
M	Number of available PN inputs	151
N_{pn}	Number of PN inputs per KC	55
A_{pn}	EPSP filter gain	1 (without sharpening)
τ_{pn}	EPSP filter decay constant	15 (without sharpening)
A_{lhi}	IPSP filter gain	0.4
τ_{lhi}	IPSP filter decay constant	10
Δ_{lhi}	LHI delay	15
K_{ggn}	GGN filter gain	380
θ_0	KC baseline threshold	5.5
τ_{ref}	Refractory period	150
bf	Biased fraction of the population	0.2 (optimal)
mh	Number of predetermined inputs to each biased KC	15 (optimal)

2.4 Acknowledgements

This work was done in collaboration with Sina Tootoonian, I did the experiments, analyses of PNs and KCs. Sina built the PN-KC model described in 2.3.6 and in Fig. 2.3. The PN-KC model also benefited from discussions with Anusha Narayan and Maria Papadopoulou. I would also like to thank Thanos Siapas and members of his lab for providing wire tetrodes for KC recordings.

CHAPTER 3

Odor Identification and Generalization in the Mushroom Body

Olfactory information is sent from the antennal lobe (AL) to two major centers in the insect brain, the mushroom body (MB) and the lateral horn (LH) (48, 53). Experiments that lesion the MB suggest that information flow through the LH is sufficient to support basic olfactory behaviors (86, 87), while the MB act as the seat of a memory trace for odors (80). Complex odor mixtures are represented by small sets of MB intrinsic neurons, the Kenyon cells (48, 53, 80). Synapses between KCs and their postsynaptic partners are likely sites of plasticity during associative learning (88). It was recently found that a form of spike-timing-dependent plasticity (STDP) exists at the synapses between KCs and a class of output neurons, known as β lobe neurons (β LN) (70). When an insect is repeatedly exposed to an odor, STDP could selectively strengthen synapses between odor driven KCs and β LN. β LN could

therefore selectively hone in onto these informative KCs (48). What is crucial then are the informational aspects of odor stimuli (e.g., category vs. identity) that are represented in the KC population response, because: (i) this representation is used by downstream neurons for associative learning and subsequent behavior output, and (ii), it has been hypothesized that the MB and LH likely extract different olfactory information from the same PN input population and an important goal must be to understand what this difference is and what it is useful for. This chapter is organized to answer four questions: (i) What types of olfactory information is extracted by KCs from the PN input? Or put it another way, what types of information is contained in the KC population response? (ii) What is the format of this representation? And how is it different to the format of the PN representation? (iii) What is this format useful for? (iv) How is the KC population response read out by downstream neurons?

3.1 Results

3.1.1 Decoding PN trajectories over time

I showed in Chapter 2 that individual KCs are odor classifiers, each making a binary decision on a subset of all possible PN inputs. What is the collective decoding capability of KCs as a population? We first examined the

sequential organization of the KC population. Fig. 3.1A shows three W-responding KCs. As described before, each responded at a particular time during the stimulus with just a few action potentials on a baseline of 0. These three KCs also responded to mixtures. Their responses, however, were not identical (Fig. 3.1A). They could vary in intensity, duration and in timing. Figure 3.1B plots the PSTHs of all KCs which responded to W (averaged over 7 trials and normalized) in grayscale, superimposed by its peak (red dots) and sorted by peak of the PSTH. As components were added, the number of mixture-responding KCs increased, in part because it included also other KCs that detected the other components of the mixture (black dots). The number of responding KCs, while always low (0.5–1% of all KCs in any 50 ms bin, Fig S2.17), thus increases with mixture size. This was a result not of increased total PN activity—which varies little with concentration and mixture complexity (\nearrow)—but rather of increased PN synchronization (Fig. S2.18). More interestingly, KC responses were distributed during and after the stimulus, with a peak within 200–300 ms after stimulus onset on average (Fig. 3.1C). The response time of these KCs have been superimposed on the corresponding PN trajectories for these stimuli (Fig. 3.1D). Because KCs decode PN trajectories piecewise in time with an integration time window of $\sim 1/2$ oscillation cycle (~ 25 ms) we conclude that all (within the limits of our limited sample) fragments of the PN trajectories are decoded by at least some KCs. Conversely, each odor (whether it is a mixture or not) is represented by a time series of sparse KC

activity vectors. The PN trajectories are thus mapped onto sparsened trajectories in KC space.

Figure 3.1 Distribution of KC response times and accuracy of odor categorization/identification from PN or KC population data.

A. KC rasters in response to W, WYZ and ALL. KCs respond reliably across trials and differ in response duration and timing.



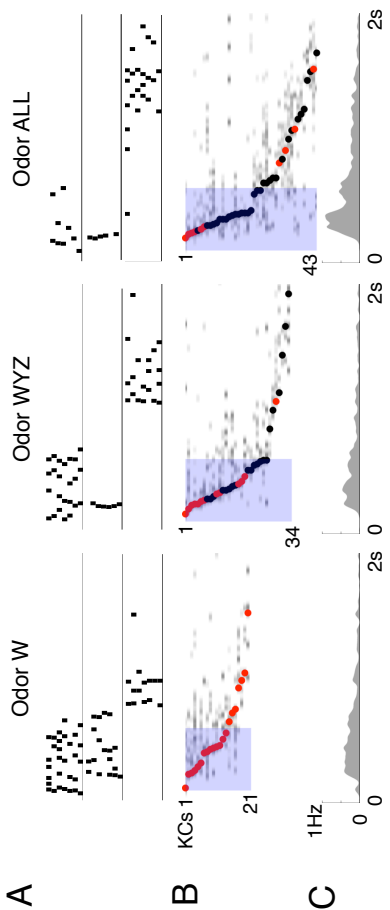


B. Activity of all recorded KCs that responded to W, WYZ, and ALL. Gray: PSTH averaged across 7 trials, normalized between 0 and 1; dots denote peak of corresponding PSTH. KCs ordered by time of peak, illustrating sequential spread of activity, especially tight within first 500 ms. Red dots: W-responding KCs; black dots: KCs that did not respond to W presented alone. Shaded bar: odor, 500 ms.

C. Instantaneous firing rates, averaged across all trials and all 209 KCs to W, WYZ, and ALL. Values are greater with larger mixtures because, on average, more KCs respond to a mixture than to individual components.

D-E. Two-trial averages of trajectories evoked by W, WYZ and ALL, analyzed over 175 PNs (50-ms consecutive bins, 4.5 s from odor onset), plotted in LLE 1-3 space. LLE computed separately for D and E (i.e., different coordinates) and plotted separately for clarity. KC PSTH peaks from **B** are superimposed on PN trajectories at corresponding times.

F. Performance of a linear classifier at decoding odor identity (dashed, chance = 2.27%) and category (dashed, chance = 50%) as functions of time (see text). Odor pulse between blue lines (500 ms). Peak of KC-decoding accuracy at black line. Shaded areas: SDs across all odor conditions. Categorization averaged over all 8 odor components (see SI Fig. 22 for breakdown).

G. Accuracy of categorization/identification as functions of number of randomly selected subsets of PNs/KCs. $x=1$ corresponds to time of black line in **F**. For PNs, decoding accuracy begins to saturate with 30% of the recorded set; accuracy never saturated with our KC set. SDs computed over 100 random subsets.



3.1.2 Decoding odor identity and category from population

KC vectors

Because KCs are read by downstream neurons in associative learning and for behavior output (48, 70, 80), we examined different aspects of- and the amount of- information preserved in the KC population vectors in short time bins. We compared the decoding of odor identity and category across the dynamic PN and KC population activity, using a linear classifier (regularized least squares classification, see Methods). Decoding accuracy was based on the responses of trials that were not included in the classifier training. Thus, the level of accuracy is what real downstream neurons could achieve in single trials by computing a weighted sum of spikes in short time windows. Decoding in this manner means that sequential patterns are lost, instead as the neuronal code (PN or KC) evolves, the linear classifier reads out one PN or KC vector at a time. Identification required attributing a particular response vector to the right odor (all-vs-all approach, 44 classes, chance = 2.27%); categorization consisted in discriminating mixtures containing a particular component from those that did not contain it (repeated for all 8 odor components, thus 8 different class partitions and averaged, see Methods; chance = 50%). Note that our PN set (175/830 \approx 21%) was proportionately much larger than for KCs (209/50,000 \approx 0.4%). The results are plotted in Fig. 3.1f as a function of time around the stimulus. As expected, identity (left) and category (right) assignment were nearly

perfect with this PN set (red). Peak performance was obtained at ~300 ms on average after stimulus onset and remained high for ~500 ms beyond odor offset. Identification and categorization success reached nearly 100% and 90%, respectively. The fact that identification has a higher success rate than categorization, even though identification is a more difficult problem ($1/44$ as opposed to $1/2$) suggests that the primary goal of AL is to optimize odor representations in PN coding space so that each odor is distinctly represented. Categorization naturally falls out of this computation, but does not appear to be a goal in itself. Remarkably, categorization success with our small KC set reached nearly 80% overall (as high as nearly 90% for individual odor components), while identification reached ~45% on average. Peak categorization performance (averaged over odors) was reached at similar times with KCs and PNs (Fig 3.1F, right). This correlation was even more striking when performance profiles were considered individually, per odor component (Fig. S3.2). The temporal profile also varied depending on the odor category, peaking as early as ~100 ms after odor onset for W, but as late as ~200 ms after odor offset for A (Fig. S3.2). Interestingly, these temporal profiles were also preserved when even small subsets of neurons were considered (Fig. S3.3–S3.9). Thus it appears that KCs as a population could extract different stimulus features at different favored times. These observations are consistent with the instantaneous, piecewise decoding of PN output by KCs (60). They also indicate that peak accuracy is not reached with uniform dynamics for all stimuli.

Because odor representations are distributed broadly across PNs and sparsely across KCs, we next quantified differences in coding capacity between PNs and KCs as a function of the number of neurons (10, 89). Using the population vectors at the times of peak KC performance (100 ms bin, vertical lines in Fig 3.1F), we progressively withdrew random fractions of each population and recomputed identification and categorization success as a function of the remaining fraction of neurons (Fig. 3.1G, see Methods). Degradation of performance at either task was significantly greater with KCs than with PNs (see fractional degradation at $x=0.5$ for example). This confirms that odor representations in the mushroom body are carried by small sets of very informative neurons, whose absence in a template-matching test has grave consequences for recognition. (Note that “proportion of all neurons” in Fig 3.1G refers to the fraction of neurons in our datasets. Plotted on an axis representing the full populations, 1—our full dataset—would correspond to 21% for PNs, and 0.4% for KCs.)

Next, we examined categorization and identification performance with and without sets of the most component-selective neurons (by ROC analysis, Fig. 2.2E–F). Our goal was to assess the sensitivity of decoding to the absence of odor-segmenting neurons. For identification and categorization (especially early in the response), KC decoding accuracy was extremely sensitive to loss of component-selective neurons. For example, identification degraded significantly if as few as the top ~3% of KCs were

removed. By contrast, PN decoding accuracy did not degrade until after the top 24% had been removed (Fig. S3.3–9). Thus, the component-selective KCs we observed experimentally (Fig. 2.2B–C) are crucial for representing both identity and category. A goal of the mushroom body thus appears to be to change odor representations towards a format where category (and identity) attributes can be retrieved from very few component selective neurons. This operation does not increase the amount of information encoded, but it greatly facilitates generalization (see below), while keeping the size of the cell assembly required for memorization small.

3.2 Discussion

3.2.1 Odor segmenting KCs leads to stimulus generalization

The discovery of odor-segmenting KCs is important for our understanding of computation in this system. Our results are illustrated in Fig 3.2A. Each row represents a KC (taken from our dataset) that expressed good segmenting properties (one KC for each one of the eight single odors). Each column represents one of the 43 stimulus conditions (paraffin oil not shown). The color of each square identifies the odor component that the corresponding KC detected. The dots indicate the responses (summed over 7 trials in 1 s and normalized) of each cell to each stimulus. For example, KC3 was a nearly perfect segmenter for citral, with only one false positive

(last of the 3-mixtures) and one false negative (8-mixture) (Note dots without a colored square indicate false positives while squares without dots indicate false negatives). The obvious features of this coding scheme are its economy and its sparsity: any odor—whether simple or composite—can be represented by a unique, low-dimensional (here, 8-D) vector of activity. The first 4- and 5-mixtures, for example, are represented by KC activity vectors that differ simply by the activation of KC6. To the limit, if every KC was a perfect detector for only one feature, then n KCs could encode $2^n - 1$ different odor feature combinations, plus baseline (0,0,...0). By contrast, a “grandmother” scheme whereby each odor is represented by a unique neuron would require $2^n - 1$ KCs to represent this many odors and mixtures. Hence, KCs implement a clever strategy. Odor representation is sparse (effective for memory formation and recall, yet not maximally sparse), but distributed such that the coding capacity for related stimuli (mixtures) is maximized.

One can argue that the coding capacity of PNs is superior to KCs, because they engage far fewer neurons (800 vs 50,000) to accomplish the same goal; indeed, the information captured by KCs is obviously present across the PN population. This is true, until one observes that PN codes are dense, and thus overlap extensively. Such a scheme is economical for encoding, but bad for storage (82, 90): any synapse modified to encode one memory will influence the encoding of other memories, thus causing

interference. Mushroom body representations are therefore the expression of a trade-off: by using odor-segmenting neurons (explicit feature-representation scheme), they maximize capacity, while minimizing total KC number. By using sparse representation vectors, they minimize overlap and interference. Both attributes are desirable for a memory system.

This mixture-encoding format has yet another advantage: it naturally leads to stimulus generalization. The logic is illustrated in Fig. 3.2B-F. Imagine an odor space of three components (A, B, C: basis vectors in that space), where each is represented by one KC (KC1-3); each odor (mixtures included) is represented by a 3-KC vector. Decoding can be achieved simply by linear separation, formalized as hyperplanes: in Fig. 3.2B, the three decoders (hyperplanes) each separate A-, B- and C- containing odors. In Fig. 3.2C, the decoders also separate AC-, AB- and BC- containing odors, allowing for multiple levels of categorization. The advantage of this ordered scheme is that it makes generalization very easy: a decoder that separates A, AC and AB (red plane, Fig 3.2B) will naturally group ABC in the same category and thus, enable generalization. This is a consequence of the fact that representation space is ordered. In Fig. 3.2D-F, the same cells and odors are plotted but each KC does not represent a basis vector. In 3.2D, the KCs have been shuffled such that all odors can still be represented by a 3-KC vector (A, B, C, AB, AC and AD are still linearly separable), but the absence of odor is now represented by spiking in KC3, and odor A by (0,0,0), both of which

make no biological sense. In Fig 3.2E, AB, AC and BC are not linearly separable anymore from, respectively, (not-AB), (not-AC) and (not-BC). A hyperplane exists that can separate B from (not-B) (example shown), but this requires reading out 3 KCs rather than just one as in the ordered scheme Fig 3.2B. In Fig 3.2F, C and (not-C) are not linearly separable (shown). Therefore, stimuli are not represented by sparse sets of randomly activated KCs (e.g., like the ASCII code). Each KC would represent some meaningful feature, and each stimulus would be encoded by the combination of relevant feature-selective KCs (82, 91). This ordered scheme allows downstream decoders of KC activity to determine not only the degree of similarity between stimuli, and but also allow for generalization of novel stimuli. Hence, the scheme we observed for mixture coding by KCs is consistent with the fulfillment of several concurrent requirements: economy of size, maximization of capacity for that size, minimization of overlap between memories, and generalization (see Table 3.1). These requirements are served by preprocessing of odor mixtures by PNs, by the architecture of PN-KC circuits and possibly, by an appropriately biased connectivity between PN and KC populations. The rules identified here for a simple olfactory system could, in principle, form the basis for the encoding of multi-dimensional signal in any sensory system with comparable requirements.





Figure 3.2. Coding principles for odor identification and generalization by KC assemblies.

A. Diagram indicating the responses of eight of our recorded component-detecting KCs (one for each of the 8 tested) to all odor conditions. Response measured as spike counts in 1-s window from odor onset summed over all 7 trials, normalized between 0 and 1 for each KC. Filled circle area represents response. Each color represents one odor. A filled circle on a colored square indicates a “true positive” (a KC recognizes the component to which it is tuned from within a mixture). An empty colored square indicates a “false negative”. Black circle alone indicates “false positive”. KC 5 detects W without mistake. In the absolute, downstream decoders need only read one KC to recognize a category: for instance, a response in KC 3 indicates the presence of odor C. Identification is possible by observing across the 8-D KC vector (e.g., discrimination between mixtures ABCD and ABCDX, vertical arrows). With perfect component-detectors, n KCs can discriminate between $2^n - 1$ odor mixtures (all combinations) and baseline. The code is both sparse and distributed (i.e., not a grand-mother scheme).

B-C. Linear classification of odors and mixtures using ordered KC encoding scheme. Schematic of KC coding space where each KC represents an odor component. With this scheme, generalization is simple: B shows hyperplanes that separate mixtures into A vs. not-A, B vs. not-B, and C vs. not-C; C shows hyperplanes that generalize mixtures into AB vs. not-AB, BC vs. not-BC, and AC vs. not-AC. Information is represented such that different kinds of generalization are easy to compute with linear classifiers.

D-F. Coding strategies with a scrambled scheme. In D, the presence of odors A, B, C, AB, AC and BC can be computed with a single linear classifier in each case, but A is represented by [0 0 0] and no-odor is represented by activation of KC3. This strategy is incompatible with experiments; KC baseline activity is 0, precluding signaling by inhibition. In E, although a hyperplane exists that separates B from not-B, odors AB/not-AB, AC/not-AC and BC/not-BC each require multiple hyperplanes for separation. In F, odors C and not-C are not separable by a single hyperplane (as shown). See text for details.

	Coding Capacity	Memory Capacity	Speed of Learning	General-ization	Interfer-ence	Fault Tolerance
Local	Very low	Limited	Very fast	None	None	None
Sparse	High	High	Fast	Good	Controlled	High
Dense	Very high	Low	Slow	Good	Strong	Very high

Table 3.1 Characteristics of different coding strategies, adapted from (82).

3.2.2 Combinatorial codes for mixtures in the MB

Populations of neurons with feature selectivity could function in two very different ways. In one view, the signals carried by the different cells are redundant and primarily independent, so that averaging over the population serves to reduce noise, as was the case with PNs. At the opposite extreme, information is transmitted by a combinatorial code, defined as the pattern of spikes and silences across the population. If different combinations of spikes and silences stand for different sensory stimuli, then averaging would discard potentially large amounts of information present only in the combinatorial code. A number of groups have presented evidence for combinatorial coding, in the cortex (92–94). To make use of a KC combinatorial code for mixtures, downstream neurons must combine inputs from different neurons in a way that distinguishes the different combinations of spiking and silence. We found that simple linear classifiers can rapidly (within 200–300 ms of odor onset) and accurately decode the category of an odor from the firing patterns of 209/50,000 KCs. Interestingly, it was recently found that a form of spike–timing–dependent plasticity (STDP) exists at the synapses between KCs and a class of output neurons, the β LNs (70). When an

insect is repeatedly exposed to odors of a particular category (e.g., AB, AC, AD), STDP could selectively strengthen synapses between A-driven KCs and β LN_s, that could allow for subsequent generalization of novel category A odors (e.g., AX, AY, AZ etc). The β LN_s also inhibit each other (70), such circuits could at least theoretically (92) also evaluate combinations of spiking and silences across sparse KC feature vectors to distinguish between highly overlapping mixtures (e.g., ABCD vs. ABCDX as in Fig. 3.2A




Fig. S3.1

Categorization of Odor Components

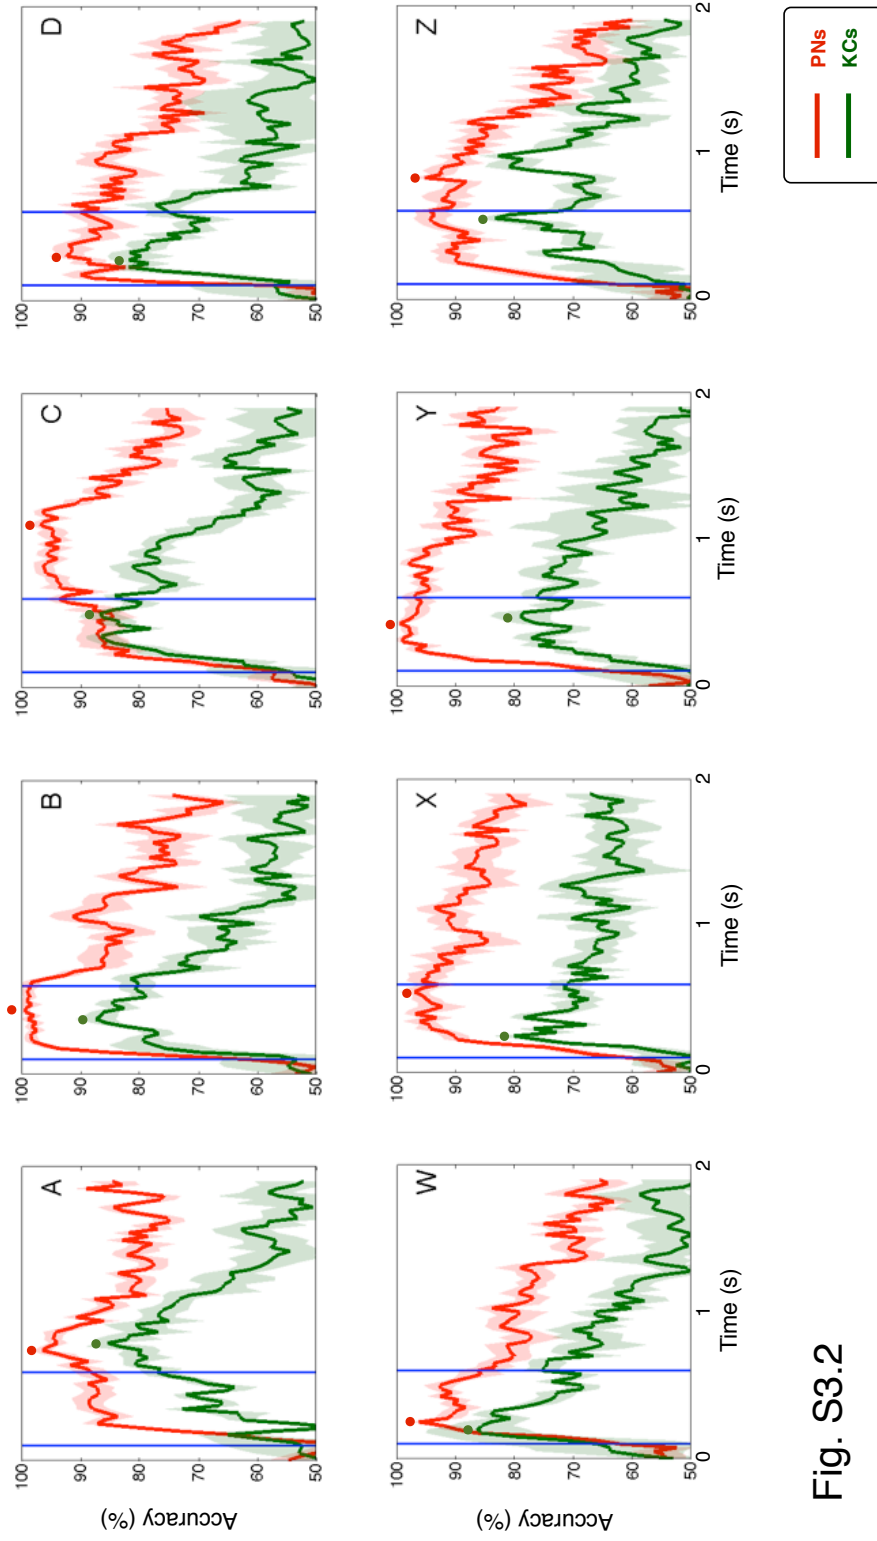


Fig. S3.2

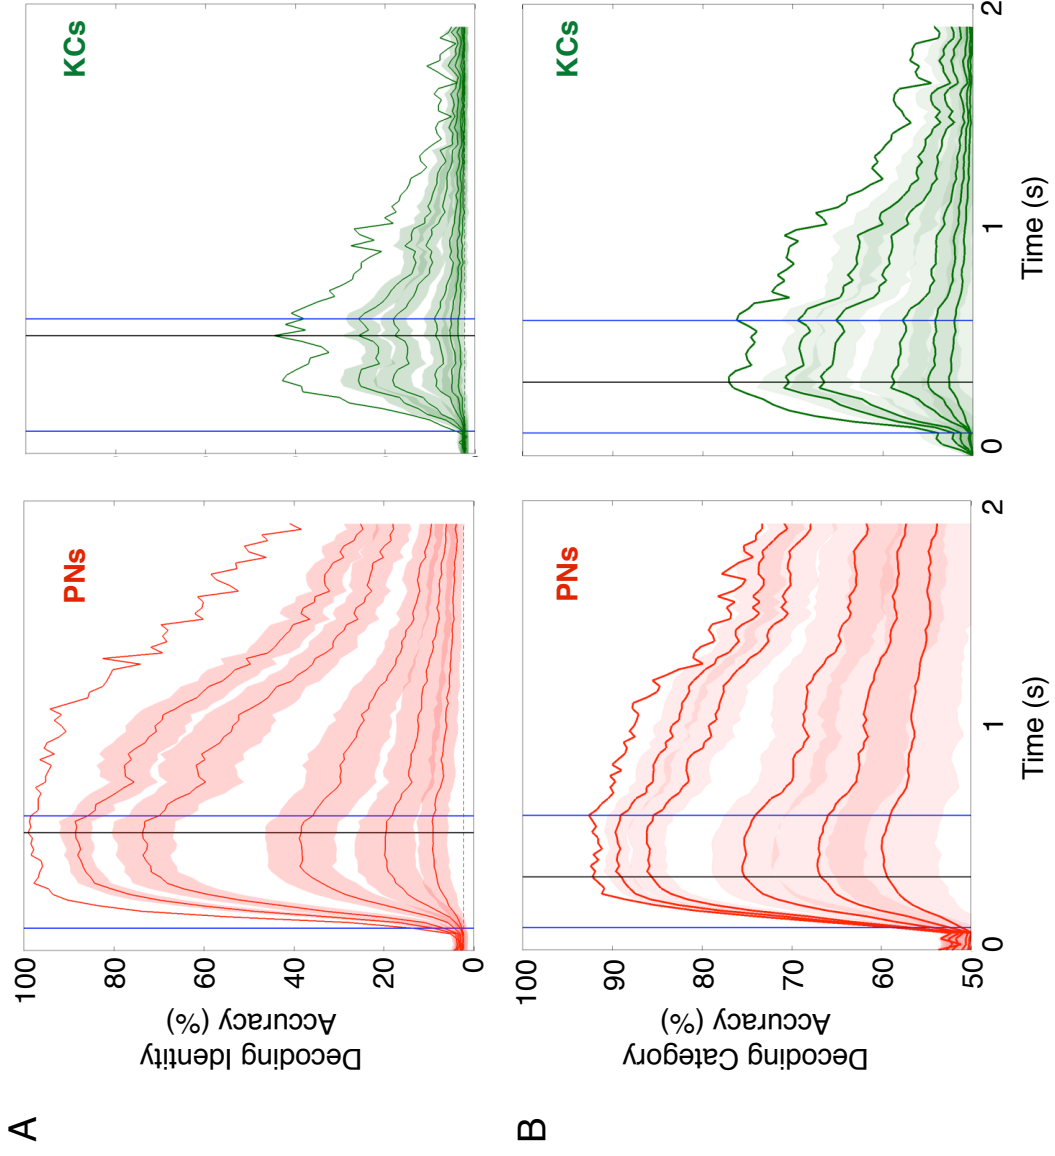


Fig. S3.3

Decoding Odor Category

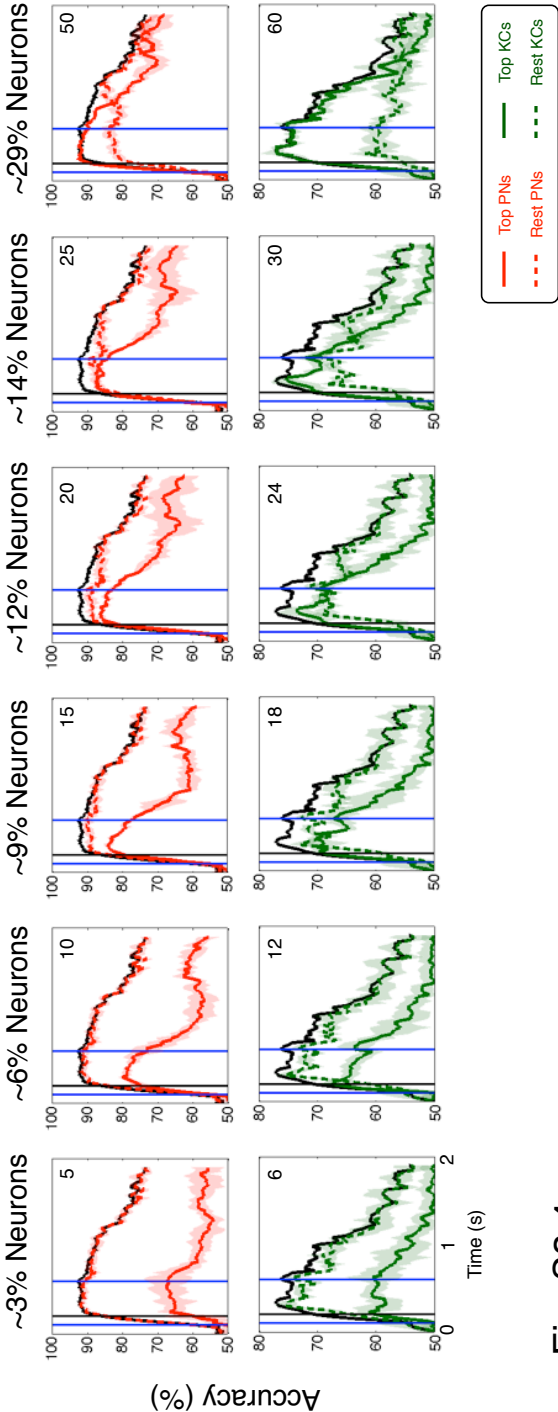


Fig. S3.4




Fig. S3.5

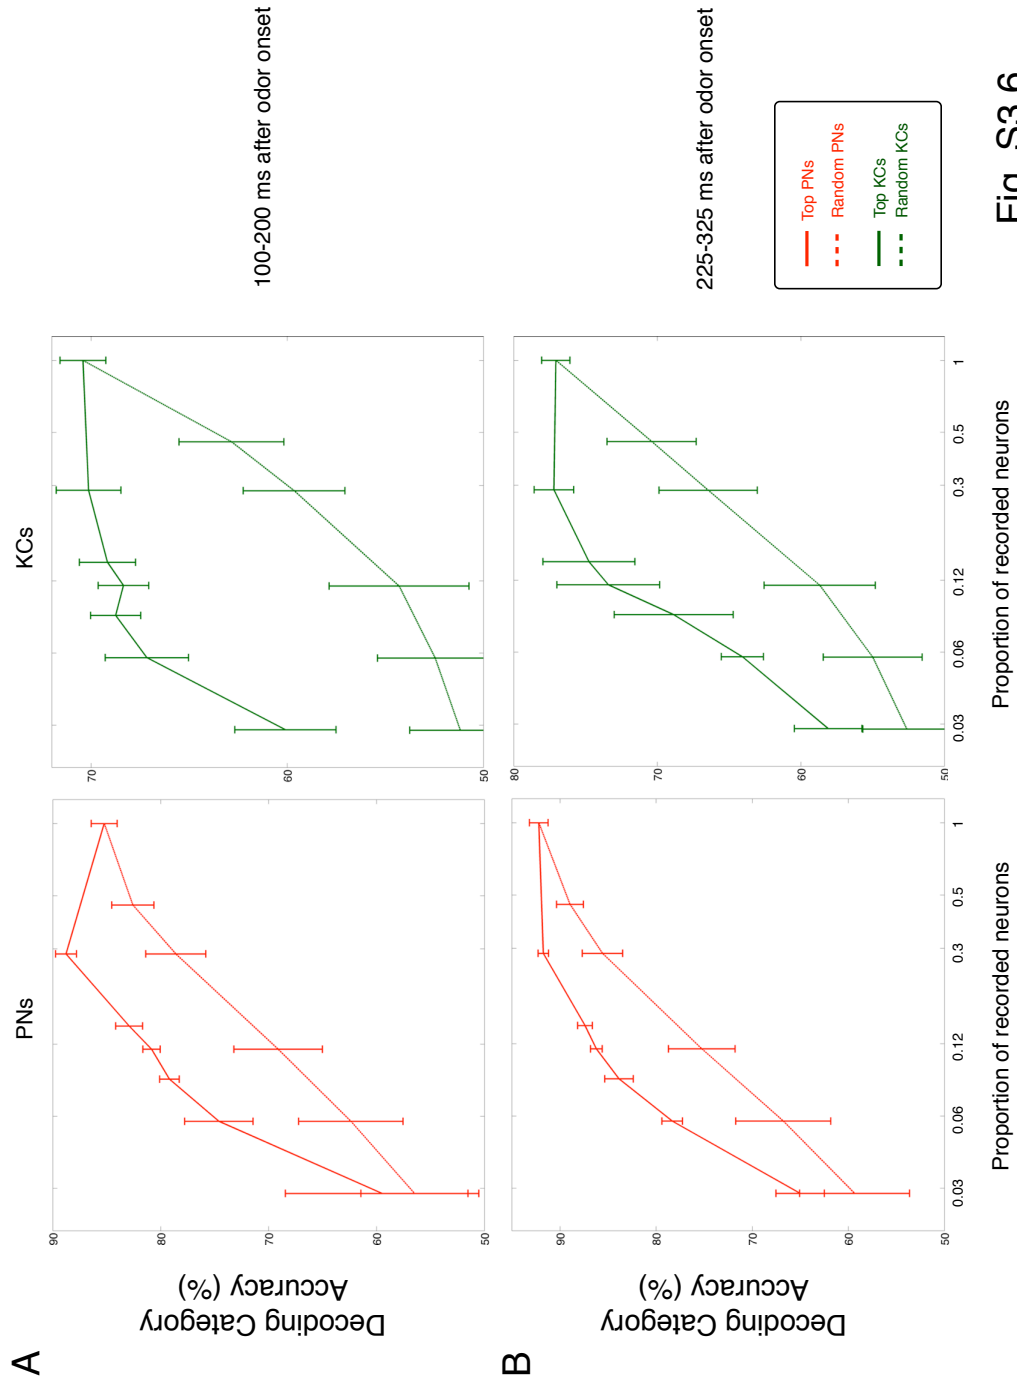


Fig. S3.6




Fig. S3.7




Fig. S3.8

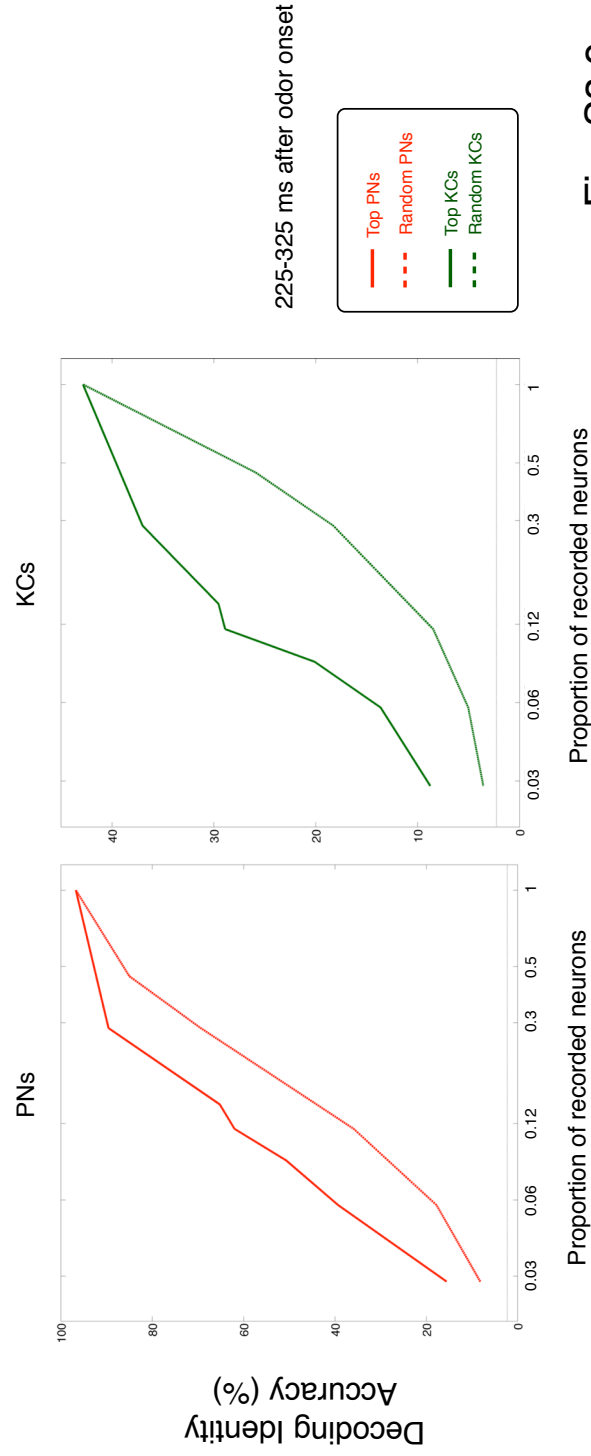


Fig. S3.9

CHAPTER 3 SUPPLEMENTARY FIGURE LEGENDS

Figure S3.1 Classification accuracy superimposed with firing rate.

A–B: Categorization and identification accuracy superimposed for comparison of time course for both PNs (**A**) and KCs (**B**). There are no obvious differences in dynamics of decoding accuracy. Instantaneous firing rate of PNs and KCs are shown in gray (spike counts in 10 ms bins convolved with 20 ms width Gaussian filter). Shaded area is SD of accuracy over different odor conditions. Notice the large variance in identification accuracy especially for KCs. Vertical blue lines indicate odor onset and offset (500 ms odor pulse). Dashed lines indicate chance performance (2.27% for identification, 50% for categorization).

Figure. S3.2 Comparison of category decoding accuracy between PNs and KCs for each single odor component.

Decoding accuracy of KCs follows faithfully that of PNs, the only exception is for odor component 'C', where the peak accuracy for KCs occurs ~500 ms before that of PNs. The timing of the peak accuracy varies across odor component categories, occurring within 100 ms of odor onset for 'W', but as long as ~200 ms after odor offset for 'A'. Thus KCs as a population could extract different stimulus features at different favored times. Average of categorization accuracy across all 8 odor components leads to Fig. 4F (right). Dots indicate time of peak performance. Shaded area represent SD over 50 bootstraps. Vertical blue lines indicate odor onset and offset (500 ms pulse). Chance performance at 50%.

Figure. S3.3 Compact and redundant coding in PNs and KCs.

A. Identification. Random subsets of 80, 50, 20, 10, 5 PNs out of a possible 175 PNs, and random subsets of 96, 60, 24, 12 and 6 out of a possible 209 KCs were used. As shown in Fig. 4G and here over time, when ~50% of PNs were removed, peak accuracy does not degrade by very much, in comparison, when ~50% of KCs were removed, peak accuracy degraded significantly. This suggests that information contained in PN ensembles is more redundant and distributed than KC ensembles. Interestingly, temporal dynamics in accuracy (e.g. double peak in KC decoding accuracy shown above)

are well preserved even for small subsets of KCs, suggesting that KC population vectors are informative at different favored times. Vertical black line indicate time of peak accuracy in both PNs and KCs. Vertical blue lines indicate odor onset and offset. Shaded area shows SD over 100 random subsets. Chance performance at 2.27%.

B. Categorization. Temporal dynamics in KC decoding accuracy (i.e., double peak) are well preserved even for small subsets of KCs, suggesting that KC population vectors are informative at different favored times. Chance performance at 50%.

Figure. S3.4 Sensitivity of *category* decoding to the absence of odor-segmenting neurons.

Decoding odor categories using the k most component-selective neurons and (all - k most component-selective) neurons. Component-selective neurons are defined as neurons with the highest AUC values by ROC analysis (see Fig. 2F and Methods). Red (PNs) and green (KCs) lines indicate performance when the top k neurons were used, the dashed lines indicate when the top k neurons were excluded and remaining neurons in the population were used (thus $(175-k)$ for PNs, and $(209-k)$ for KCs). Solid black lines indicates decoding accuracy using the full ensemble (175 PNs or 209 KCs). The vertical black line denote 100 ms after odor onset, during this early part of the response, 18/209 KCs (~9%) achieves an accuracy that is equal to the entire population, removing these 18 KCs, with $(209-18)$ remaining KCs, performance degrades substantially. In comparison, 15/175 PNs (~9%) also achieves performance equal to the full population, however, even when these PNs were excluded, $(175-15)$ remaining PNs also perform equal to the full population. This was also the case when the 50 (~29%) most component-selective PNs was removed, the remaining $(175-50)$ PNs still perform well above chance, illustrating redundant information in the PN population. In contrast, removing the most component-selective KCs degraded performance significantly, especially in the early part of the response. k is indicated in the upper right hand in each box plot. Each column correspond to equivalent proportion of recorded PNs and KCs for given values of k , e.g., $k = 5$ for PNs and $k = 6$ for KCs are both ~3% of the full recorded ensemble; $k = 10$ for PNs and $k = 12$ for KCs are both ~6% of the full recorded ensemble and so forth. Vertical blue lines indicate odor onset and offset. Shaded area shows SD over 50 bootstraps. Chance performance at 50%.

Figure. S3.5 Category decoding of KCs is more sensitive to absence of odor-segmenting neurons.

Decoding odor categories using the k most component-selective neurons as a proportion of all recorded neurons for two different time periods during stimulus response. The solid line represents decoding accuracy using the k most most-selective neurons, the dashed line at the corresponding x-axis points show decoding accuracy excluding the k most-selective neurons. The vertical black line indicates the proportional point when accuracy using the top k neurons equals accuracy using all- k neurons.

A. For KCs, accuracy was extremely sensitive to removal of component-selective neurons, and performance by the top ~6% did better than the remaining ~94% of KCs. In comparison, PN decoding accuracy were not effected until the top ~14% PNs were removed, even when the top ~29% were removed, accuracy was still well above chance.

B. Differences between PN and KC decoding observed in **(A)** not as stark at a later time point, but still significant. Error bars show SD over 50 bootstraps.

Figure. S3.6 Comparison of category decoding accuracy using k most component-selective and k randomly selected neurons.

A-B: Decoding odor categories using the k most component-selective neurons as a proportion of all recorded neurons for two different time periods during stimulus response. The solid line represents decoding accuracy using the k most most-selective neurons, the dashed line at the corresponding x-axis points show decoding accuracy using k random subsets. For a given proportion of recorded neurons, the top k neurons perform better than random subsets, but this difference is much larger for KCs than PNs, especially in the early period after odor onset **(A)**. Error bars show SD over 50 bootstraps.

Figure. S3.7 Sensitivity of *identity* decoding to the absence of odor-segmenting neurons.

Decoding odor identity using the k most component-selective neurons. Component-selective neurons are defined as neurons with the highest AUC values by ROC analysis

(see Fig. 2F and Methods). Red and green lines indicate performance when the top k neurons were used, the dashed lines indicate when the top k neurons were excluded and remaining neurons in the population were used (thus $(175-k)$ for PNs, and $(209-k)$ for KCs). The black line indicates when the full population (175 PNs or 209 KCs) was used. The vertical black line indicates the first peak in KC decoding accuracy. KC decoding accuracy is extremely sensitive to loss of component-selective KCs, performance degrades significantly even when just 6 KCs were removed. 60 most component-selective KCs contains nearly all identity information. When these 60 KCs were removed, performance degraded from ~45% to ~10%. In comparison, removing component-selective PNs has little influence on overall accuracy, e.g., even when 50 PNs (~29%) were removed, PN decoding accuracy degraded only marginally (from ~100% to ~80%). We conclude that component-selective neurons play a fundamental role in identity coding in KCs, they are not nearly as crucial for identity coding in PNs. k is indicated in the upper right hand in each box plot. Each column corresponds to equivalent proportion of recorded PNs and KCs for given values of k . Vertical blue lines indicate odor onset and offset. Shaded area shows SD over all odor conditions. Chance performance at 2.27%.

Figure. S3.8 Identity decoding of KCs is more sensitive to absence of odor-segmenting neurons.

Decoding odor categories using the k most component-selective neurons as a proportion of all recorded neurons. The solid line represents decoding accuracy using the k most most-selective neurons, the dashed line at the corresponding x-axis points show decoding accuracy excluding the k most-selective neurons. The vertical black line indicates the proportional point when accuracy using the top k neurons equals accuracy using all- k neurons. For KCs, accuracy was extremely sensitive to removal of component-selective neurons, accuracy degrades as soon as any KCs are removed. PN decoding accuracy is not affected until about ~29% of neurons were removed.

Figure. S3.9 Comparison of identity decoding accuracy using k most component-selective and k randomly selected neurons.

Decoding odor categories using the k most component-selective neurons as a proportion of all recorded neurons. The solid line represents decoding accuracy using

the k most most-selective neurons, the dashed line at the corresponding x-axis points show decoding accuracy using random subsets. For a given proportion of recorded neurons, the top k neurons perform better than random subsets, but this difference is much larger for KCs than PNs.

3.3 Methods

3.3.1 Mean KC spike latency

To quantify the preservation of temporal structure across the KC population to different odor conditions, we computed for each KC a measure of mean spike latency. For each KC, PSTHs were computed using a 20ms Gaussian smoothing kernel, averaged across 7 trials and baseline subtracted. Mean latency was defined as the peak of the PSTH.

3.3.2 Population decoding

Regularized least squares classification (RLSC)

To estimate the information carried by PN and KC ensembles about odor component and identity in single trials, we used a decoding based approach (10, 18). The linear classifier was provided with spike counts in 4 consecutive bins (25 msec each bin) across all PNs (175) and KCs (209) and computed over 2 s shortly after odor onset. The classifier consisted of a weighted sum of PN or KC inputs. The weights were estimated using

regularized least squares (95). This approach can be thought of as multiple linear regression with a constant term. Multiple linear regression cannot determine the weights unambiguously if the sample matrix is ill conditioned, this is often the case with few trials or few spikes (as was the case with KCs). The formulation for RLSC is below:

$$w = (X^T X + \lambda I)^{-1} X^T y \quad \text{Equation (3.1)}$$

The $T \times n$ matrix X contains spike counts across all cells, each row is one trial (T contains $T/2$ positive trials, and $T/2$ negative trials), each column represent spike counts in one cell (n columns represents n cells). w is the $n \times 1$ weight vector, a unique weight for each cell. y is the $T \times 1$ vector of class labels (+1 and -1). I is the $n \times n$ identity matrix, and λ is the scalar regularizer. The larger λ is, the more constraints are placed on the solution, the smaller λ is, the closer the solution is to multiple linear regression. Even a small value of the regularizer punishes unrealistically large weights and ensures stability. Regularization becomes particularly important when the number of input variables (neurons) outnumber the number of training examples, as was the case here. There is an optimal value for λ , in this thesis values of 0.01, 0.1, and 1 were tried and no significant differences were observed. Therefore λ was kept constant at 1 throughout. The number trials in each class during training were always kept the same to avoid any decoding bias. Where the number of positive and negative trials were different, we repeated k (20 or 50 bootstraps)

random sampling to equalize the number of positive and negative trials. The decoding accuracy of the classifier was estimated using leave-one-out cross validation for all training samples available.

Decoding odor identity

To decode odor identity (i.e., which of the 44 odor conditions were presented), we used all-vs-all multiclass decoding, one time bin at a time. The number of spikes was counted in each 100 msec time bin sampled at 25 msec intervals with data from each time bin being classified independently, leading to a slight temporal smoothing. We built 44x43 binary classifiers (e.g., A vs Ahigh, A vs B, A vs C, A vs D, A vs AB, ... Z vs A, Z vs AB etc) for every time bin and every trial using all trials minus one, the leave-one-out trial was classified using all binary classifiers and assigned the class with the maximum votes across classifiers. In cases where there was a tie, the leave-one-out trial was assigned randomly to one of the leading vote getters. Chance performance was 1/44 or 2.27%.

Decoding odor component

To decode odor component or class information (as we define it here), we built 8 different classifiers (A vs Not-A, B vs Not-B, C vs Not-C, D vs Not-D, W vs Not-W, X vs Not-X, Y vs Not-Y, Z vs Not-Z), one for each odor component. Because odor mixtures contain many different components, an odor mixture could be classified as belonging to multiple classes. To decode only the odor component information, careful consideration must be applied to how the two classes are partitioned. First, we must partition the two classes to decode only the odor component; that is, if we want to decide between A and Not-A, and if ABC is included in the positive class and WXY in the negative class, then we might be decoding the difference between ABC and WXY, and *NOT* that between A and Not-A. A more appropriate partitioning might be all odor conditions containing A in one class (e.g., A, AB, ABC, AC) and the exact same mixtures minus A (e.g., B, C, BC). In the latter case, if a difference is detected between the two classes, it must have been because of the presence/absence of A. We thus tried to partition odor conditions in such a way as to match this paradigm. Secondly, there is a positive correlation between the number of KCs activated and the number of odor components in the mixture n (see Fig. S2.17); thus we must also match the two classes for number of components; indeed, it would be inappropriate to have A, AB, AC and AD in one class and BCWX, BDWX, DWYZ etc in the other. In such a scheme, we could be decoding for <number of components in the mixture> and *NOT* the <component

identity>. We thus matched the two classes to have approximately the same number of n - mixtures. The class partitions for each classifier are given below. We performed k bootstraps (20 or 50) of all available trials, where in each bootstrap, the number of positive and negative trials were kept equal. Because this classification is binary, chance performance was at 50%.

Odor A vs. Not A

A: A (4x), A, AB, AC, AD, AX, ABC, ACD, AXZ, ABCD, ABCX, ABCDX, ABCWX, ADWYZ, AWXYZ

A': C (4x), B, C, D, X, BC, DW, XZ, BCX, BDW, DXY, WXY, WYZ, BCWX, BDWX, DWYZ, WXYZ, BCWXZ

Odor B vs. Not B

B: B, AB, BC, ABC, BCX, BDW, ABCD, ABCX, BCWX, BDWX, ABCDX, ABCWX, BCWXZ

B': A (4x), A, C (4x), C, X, AC, AD, AX, DW, WX, ACD, AXZ, DXY, WXY, DWYZ, WXYZ, ADWYZ, AWXYZ

Odor C vs. Not C

C: C (4x), C, AC, BC, ABC, ACD, BCX, ABCD, ABCX, BCWX, ABCDX, ABCWX, BCWXZ

C': A (4x), W (4x), A, B, X, AB, AD, WX, AXZ, BDW, DXY, BDWX, DWYZ, WXYZ, ADWYZ, AWXYZ

Odor D vs. Not D

D: D, AD, DW, ACD, BDW, DXY, ABCD, BDWX, DWYZ, ABCDX, ADWYZ

D': W (4x), A, B, W, AC, XY, ABC, WXY, WYZ, ABCX, BCWX, WXYZ, ABCWX, AWXYZ, BCWXZ

Odor W vs. Not W

W: W (4x), W, DW, WX, WY, WZ, BDW, WXY, WYZ, BCWX, BDWX, DWYZ, WXYZ, ABCWX, BCWXZ

W': A (4x), C (4x), B, D, X, Y, Z, BC, XY, XZ, YZ, ABC, ACD, AXZ, BCX, DXY, ABCD, ABCX, ABCDX

Odor X vs. Not X

X: X, AX, WX, XY, XZ, AXZ, BCX, DXY, WXY, ABCX, BCWX, BDWX, WXYZ, ABCDX, AWXYZ

X': A (4x), C (4x), W (4x), A, D, W, Y, Z, BC, WY, WZ, YZ, ABC, ACD, BDW, WYZ, ABCD, DWYZ, ADWYZ

Odor Y vs. Not Y

Y: Y, WY, XY, YZ, DXY, WXY, WYZ, DWYZ, WXYZ, ADWYZ, AWXYZ

Y': D, W, X, Z, DW, WX, WZ, XZ, AXZ, BDW, BCWX, BDWX, ABCWX, BCWXZ

Odor Z vs. Not Z

Z: Z, WZ, XZ, YZ, AXZ, WYZ, DWYZ, WXYZ, ADWYZ, AWXYZ, BCWXZ

Z': W, X, Y, AX, DW, XY, WY, BDW, WXY, DXY, BCWX, BDWX, ABCDX, ABCWX

Summary and Conclusions

I will now conclude by briefly recapitulating the main results of my PhD work, discuss their overall significance and point out open questions that remain to be answered.

4.1 An elegant solution for odor recognition and generalization

In previous electrophysiological and imaging studies of KCs in locusts, *Drosophila*, and moths, they have been found to be sparse (53, 61); selective for specific stimulus history (71); sometimes invariant to changes in concentration (7); likely built from maximally separate populations of projection neurons (59); and even by recording from similar KCs using genetic labels, they fail to show consistent odor tuning from fly to fly (68). Taken together, these results point to KCs as highly selective responders for very specific combinations of odors that arises from random connectivity of PN inputs. The sparseness of KCs are believed to facilitate synthetic odor representations that lack detail of

components, which is in fact how odors are generally perceived (multi-component mixtures that excite a wide array of receptors are nonetheless perceived as a single synthetic odor, e.g., coffee). The significance of this thesis is 3-fold: first, the discovery of component-selective Kenyon cells was surprising to us because KCs were expected to be selective for mixtures and not invariant to them. Individually, they are explicit representations of odor features (i.e., they explicitly convey odor category) and in a small population can represent odor identity with high accuracy. Second, the connectivity between PNs to KCs were expected to be random because randomness maximizes selectivity (59), the results of the PN-KC model fed with experimental PN data, indicate that randomness (given the biological constraints of this circuit) cannot account for the distribution of component-selective KCs observed experimentally. Therefore, PN-KC connectivity is either deterministic¹ or PN-KC synapses can be altered through learning. And third, due to very sparse firing of KCs, and given that downstream decoders looks across many of these KCs at once (~5,000 KCs), it was not expected that both category and identity information can be retrieved from so few KCs (<0.4% of entire population) at high accuracy. This is the first instance that a decoding method has been applied to KC population data, and the

¹ PN-KC connectivity are unlikely to be completely random if not for the simple reason that random biological circuits have never been discovered in the nervous system (although its not clear how randomness can be proved experimentally).

first time that both identity and category information has been decoded from KCs. Taken together, these results show that individual KCs are odor classifiers (or feature detectors). They sample $50 \pm 15\%$ of all PNs in a deterministic fashion (either via genetic encoding or learning). Many KCs are explicit representation of odor category, and the combination of just a few KC classifiers enables both identity and category information to be read out by a linear classifier at high accuracy, and more importantly allows for both generalization across mixture conditions, and efficient memory storage by virtue of their sparseness.

4.2 Open questions

The results of this thesis have advanced our understanding of odor representations in the insect olfactory system, especially that of the MB. However several open questions remain:

(i) What types of deterministic connectivity will give rise to component-selective KCs? If learning is involved, what form of learning is it, and how does it operate? How are KCs involved in learning associations?

(ii) What is the limit to generalization? Presumably KCs will not keep respond continuously to increasing n components². However, the limit is likely to be much greater than 8 components and maybe dependent on odor components.

(iii) Demonstrating that identity and category information are present across populations of KCs does not mean that there are used by downstream decoders. Given that (1) there are limited downstream neurons that read the KC population: ~60 β LN, probably on the same order for the α lobe neurons, (in flies, there is also the γ lobe), and (2) that these synapses are plastic, information from KCs must be read in a very selective way and under the influence of learning. One logical next step would be to examine whether odor identity and category information could be decoded from β LN under various scenarios.

(iv) One speculation resulting from the work in this thesis is that there maybe a bias towards encoding categorical information across KCs. If this assumption is true, then there will be a difference in the odor information arriving in the LH, one copy about identity directly from PNs, one copy about category from KCs via MB extrinsic neurons, such a scheme would suggest that LH neurons would be computing a difference.

² This can not occur because GGN operates to set an upper bound on KC activity ratio.

What are the LH neurons doing? This area has almost uncharted and should be investigated.

BIBLIOGRAPHY

1. D. Marr, L. Vaina, *Proc R Soc Lond B Biol Sci* **214**, 501 (Mar 22, 1982).
2. J. J. DiCarlo, D. D. Cox, *Trends Cogn Sci* **11**, 333 (Aug, 2007).
3. S. Ullman, *MIT Press*, (1996).
4. E. T. Rolls, *Neuron* **27**, 205 (Aug, 2000).
5. C. D. D. DiCarlo J., *Trends in Cognitive Sciences* **11**, 333 (2007).
6. S. Ullman, S. Soloviev, *Neural Netw* **12**, 1021 (Oct, 1999).
7. M. Stopfer, V. Jayaraman, G. Laurent, *Neuron* **39**, 991 (Sep 11, 2003).
8. B. A. Olshausen, C. H. Anderson, D. C. Van Essen, *J Neurosci* **13**, 4700 (Nov, 1993).
9. Z. Kourtzi, J. J. DiCarlo, *Curr Opin Neurobiol* **16**, 152 (Apr, 2006).
10. C. P. Hung, G. Kreiman, T. Poggio, J. J. DiCarlo, *Science* **310**, 863 (Nov 4, 2005).
11. R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, I. Fried, *Nature* **435**, 1102 (Jun 23, 2005).
12. G. Laurent, *Nat Rev Neurosci* **3**, 884 (Nov, 2002).
13. J. J. Hopfield, *Proc Natl Acad Sci U S A* **96**, 12506 (Oct 26, 1999).
14. A. Dittman, T. Quinn, *J Exp Biol* **199**, 83 (1996).
15. D. G. Wallace, B. Gorny, I. Q. Whishaw, *Behav Brain Res* **131**, 185 (Apr 1, 2002).
16. K. O. Johnson, *J Neurophysiol* **43**, 1793 (Jun, 1980).
17. K. O. Johnson, *J Neurophysiol* **43**, 1771 (Jun, 1980).
18. E. M. Meyers, D. J. Freedman, G. Kreiman, E. K. Miller, T. Poggio, *J Neurophysiol* **100**, 1407 (Sep, 2008).
19. S. a. K. Theodoridis, K., *Pattern Recognition*. (Elsevier, ed. 3, 2006).
20. C. Y. Su, K. Menuz, J. R. Carlson, *Cell* **139**, 45 (Oct 2, 2009).
21. R. Haddad *et al.*, *Nat Methods* **5**, 425 (May, 2008).
22. R. M. Khan *et al.*, *J Neurosci* **27**, 10015 (Sep 12, 2007).
23. W. S. Cain, M. Drexler, *Ann N Y Acad Sci* **237**, 427 (Sep 27, 1974).
24. D. G. Laing, G. W. Francis, *Physiol Behav* **46**, 809 (Nov, 1989).
25. A. Livermore, D. G. Laing, *J Exp Psychol Hum Percept Perform* **22**, 267 (Apr, 1996).
26. A. Livermore, D. G. Laing, *Physiol Behav* **65**, 311 (Nov 15, 1998).
27. A. Livermore, D. G. Laing, *Percept Psychophys* **60**, 650 (May, 1998).
28. M. D. Rabin, W. S. Cain, *J Exp Psychol Learn Mem Cogn* **10**, 316 (Apr, 1984).
29. X. C. Lu, B. M. Slotnick, *Neuroscience* **84**, 849 (Jun, 1998).
30. N. Uchida, Z. F. Mainen, *Nat Neurosci* **6**, 1224 (Nov, 2003).
31. B. D. Rubin, L. C. Katz, *Neuron* **23**, 499 (Jul, 1999).

32. R. I. Wilson, Z. F. Mainen, *Annu Rev Neurosci* **29**, 163 (2006).
33. A. Jinks, D. G. Laing, *Perception* **28**, 395 (1999).
34. C. D. Brody, J. J. Hopfield, *Neuron* **37**, 843 (Mar 6, 2003).
35. O. Hendin, D. Horn, J. J. Hopfield, *Proc Natl Acad Sci U S A* **91**, 5942 (Jun 21, 1994).
36. R. N. Shepard, *Science* **237**, 1317 (Sep 11, 1987).
37. G. A. Wright, M. Carlton, B. H. Smith, *Behav Neurosci* **123**, 36 (Feb, 2009).
38. G. A. Wright, J. A. Mustard, S. M. Kottcamp, B. H. Smith, *J Exp Biol* **210**, 4024 (Nov, 2007).
39. G. A. Wright, A. Lutmerding, N. Dudareva, B. H. Smith, *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* **191**, 105 (Feb, 2005).
40. G. A. Wright, B. D. Skinner, B. H. Smith, *J Chem Ecol* **28**, 721 (Apr, 2002).
41. G. A. Wright, S. M. Kottcamp, M. G. Thomson, *PLoS One* **3**, e1704 (2008).
42. G. Laurent, M. Naraghi, *J Neurosci* **14**, 2993 (May, 1994).
43. K. MacLeod, A. Backer, G. Laurent, *Nature* **395**, 693 (Oct 15, 1998).
44. E. A. Hallem, J. R. Carlson, *Cell* **125**, 143 (Apr 7, 2006).
45. E. A. Hallem, A. Nicole Fox, L. J. Zwiebel, J. R. Carlson, *Nature* **427**, 212 (Jan 15, 2004).
46. G. S. Suh *et al.*, *Nature* **431**, 854 (Oct 14, 2004).
47. Y. Shang, A. Claridge-Chang, L. Sjulson, M. Pypaert, G. Miesenbock, *Cell* **128**, 601 (Feb 9, 2007).
48. N. Y. Masse, G. C. Turner, G. S. Jefferis, *Curr Biol* **19**, R700 (Aug 25, 2009).
49. T. Faber, J. Joerges, R. Menzel, *Nat Neurosci* **2**, 74 (Jan, 1999).
50. M. Bazhenov *et al.*, *Neuron* **30**, 569 (May, 2001).
51. M. Bazhenov *et al.*, *Neuron* **30**, 553 (May, 2001).
52. M. Wehr, G. Laurent, *Nature* **384**, 162 (Nov 14, 1996).
53. J. Perez-Orive *et al.*, *Science* **297**, 359 (Jul 19, 2002).
54. G. Laurent *et al.*, *Annu Rev Neurosci* **24**, 263 (2001).
55. N. J. Strausfeld, L. Hansen, Y. Li, R. S. Gomez, K. Ito, *Learn Mem* **5**, 11 (May-Jun, 1998).
56. R. A. Campbell, G. C. Turner, *Curr Biol* **20**, R11 (Jan 12, 2010).
57. M. Mizunami, J. M. Weibrecht, N. J. Strausfeld, *J Comp Neurol* **402**, 520 (Dec 28, 1998).
58. O. Mazor, G. Laurent, *Neuron* **48**, 661 (Nov 23, 2005).
59. R. A. Jortner, S. S. Farivar, G. Laurent, *J Neurosci* **27**, 1659 (Feb 14, 2007).
60. J. Perez-Orive, M. Bazhenov, G. Laurent, *J Neurosci* **24**, 6037 (Jun 30, 2004).
61. G. C. Turner, M. Bazhenov, G. Laurent, *J Neurophysiol* **99**, 734 (Feb, 2008).

62. B. A. Olshausen, D. J. Field, *Curr Opin Neurobiol* **14**, 481 (Aug, 2004).
63. R. H. Hahnloser, A. A. Kozhevnikov, M. S. Fee, *Nature* **419**, 65 (Sep 5, 2002).
64. C. Poo, J. S. Isaacson, *Neuron* **62**, 850 (Jun 25, 2009).
65. G. S. Suh *et al.*, *Curr Biol* **17**, 905 (May 15, 2007).
66. A. M. Wong, J. W. Wang, R. Axel, *Cell* **109**, 229 (Apr 19, 2002).
67. G. S. Jefferis *et al.*, *Cell* **128**, 1187 (Mar 23, 2007).
68. M. Murthy, I. Fiete, G. Laurent, *Neuron* **59**, 1009 (Sep 25, 2008).
69. heisenberg, *Nat Rev Neurosci*, (2003).
70. S. Cassenaer, G. Laurent, *Nature* **448**, 709 (Aug 9, 2007).
71. B. M. Broome, V. Jayaraman, G. Laurent, *Neuron* **51**, 467 (Aug 17, 2006).
72. G. Laurent, H. Davidowitz, *Science* **265**, 1872 (Sep 23, 1994).
73. C. Linster, B. A. Johnson, A. Morse, E. Yue, M. Leon, *J Neurosci* **22**, 6842 (Aug 15, 2002).
74. N. M. Abraham *et al.*, *Neuron* **44**, 865 (Dec 2, 2004).
75. N. Uchida, Z. F. Mainen, *Front Syst Neurosci* **1**, 3 (2007).
76. S. Bhagavan, B. H. Smith, *Physiol Behav* **61**, 107 (Jan, 1997).
77. Z. F. Mainen, *Curr Opin Neurobiol* **16**, 429 (Aug, 2006).
78. J. L. Hurst, R. J. Beynon, *Bioessays* **26**, 1288 (Dec, 2004).
79. S. L. Brown, J. Joseph, M. Stopfer, *Nat Neurosci* **8**, 1568 (Nov, 2005).
80. M. Heisenberg, *Nat Rev Neurosci* **4**, 266 (Apr, 2003).
81. T. Fawcett, *Pattern Recogn Lett* **27**, 861 (Jun, 2006).
82. P. Foldiak, in *The Handbook of Brain Theory and Neural Networks*, M. A. Arbib, Ed. (MIT Press, 2002).
83. P. A. Rhodes, *Neural Comput* **20**, 2000 (Aug, 2008).
84. C. Pouzat, O. Mazor, G. Laurent, *J Neurosci Methods* **122**, 43 (Dec 31, 2002).
85. S. T. Roweis, L. K. Saul, *Science* **290**, 2323 (Dec 22, 2000).
86. J. S. de Belle, M. Heisenberg, *Science* **263**, 692 (Feb 4, 1994).
87. J. B. Connolly *et al.*, *Science* **274**, 2104 (Dec 20, 1996).
88. A. C. Keene, S. Waddell, *Nat Rev Neurosci* **8**, 341 (May, 2007).
89. L. F. Abbott, E. T. Rolls, M. J. Tovee, *Cereb Cortex* **6**, 498 (May–Jun, 1996).
90. D. J. Field, *Neural Computation* **6**, 559 (Jul, 1994).
91. H. B. Barlow, *Perception* **1**, 371 (1972).
92. P. J. schneidman E, Harris RA, Bialek W, Berry MJ 2nd, *q-bio.NC/0607017*, (2006).
93. T. J. Gawne, B. J. Richmond, *J Neurosci* **13**, 2758 (Jul, 1993).
94. S. Panzeri, S. R. Schultz, A. Treves, E. T. Rolls, *Proc Biol Sci* **266**, 1001 (May 22, 1999).
95. R. Rifkin, Yeo, G., and Poggio, T., in *Advances in Learning Theory: Methods, Model and Applications*, H. Suykens, Basu, Micchelli and

Vandewalle, Ed. (VIOS Press, Amsterdam, 2003), vol. Vol. 190, pp. 131-154.