

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Mechanisms for optimal decision making in small neural circuits /

Permalink

<https://escholarship.org/uc/item/9z64f7f4>

Author

Calhoun, Adam J.

Publication Date

2014

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Mechanisms for optimal decision making in small neural circuits

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Neurosciences (specialization in Computational Neuroscience)

by

Adam J. Calhoun

Committee in charge:

Professor Tatyana Sharpee, Chair
Professor Terrence Sejnowski, Co-Chair
Professor Sreekanth Chalasani
Professor William Kristan
Professor Charles Stevens
Professor Jing Wang

2014

Copyright

Adam J. Calhoun, 2014

All rights reserved.

The dissertation of Adam J. Calhoun is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Chair

University of California, San Diego

2014

DEDICATION

To family and friends, but mostly to Priya

EPIGRAPH

“The greatest happiness is to know the source of unhappiness.”

“It seems, in fact, as though the second half of a man’s life is made up of nothing, but the habits he has accumulated during the first half.”

Fyodor Dostoevsky

TABLE OF CONTENTS

Signature Page	iv
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures and Tables	viii
Acknowledgments	ix
Vita	x
Abstract of the Dissertation	xi
1. Introduction	1
1.1 Circuits to extract information	1
1.2 Biology of <i>C. elegans</i>	7
1.3 Optimal behavior	10
1.4 References	14
2. Neural mechanisms for evaluating uncertainty in <i>Caenorhabditis elegans</i>	22
2.1 Introduction	22
2.2 Results	24
2.2.1 Prior food experience modifies local search behavior	24
2.2.2 Environmental variability can quantitatively predict local search	26
2.2.3 ASI and ASK sensory neurons are specialized for evaluating variance	32
2.2.4 Dopamine suppresses local search behavior	37
2.2.5 CREB influences acquisition time	42
2.3 Discussion	45
2.3.1 Environmental variance modifies local search	47
2.3.2 Dopamine and CREB signaling interpret sensory information	48
2.3.3 Two circuit motifs estimate uncertainty	50
2.4 Experimental procedures	51
2.4.1 Strains and transgenics	51
2.4.2 Learning assay	52
2.4.3 Calcium imaging	53
2.4.4 Imaging bacteria patches	54
2.4.5 Simulations	55
2.4.6 Prediction and Maximum Noise Entropy	55
2.5 Acknowledgments	60
2.6 References	61
3. Emergent discrete decisions within continuous maximally informative search strategies	66
3.1 Abstract	66
3.2 Introduction	67
3.3 Results	70
3.4 Discussion	76
3.5 Materials and methods	81
3.5.1 Quantification of animal behavior	81
3.5.2 Computation of infotaxis trajectories	82

3.5.3 Setting parameters of phenomenological drift-diffusion model	82
3.6 Acknowledgments.....	82
3.7 References.....	83
4. Conclusion	86

LIST OF FIGURES AND TABLES

Figure 2.1. On food experience modifies off-food behavior	25
Figure 2.2. Worms on differently-sized food patches experience different sensory environments.....	27
Figure 2.3. Specifics of learning.....	28
Figure 2.4. Measuring bacterial gradients.....	30
Figure 2.5. Quantifying on-food behavior	31
Figure 2.6. A filter predicts the number of turns during local search.....	33
Figure 2.7. Learning requires a unique network.....	34
Figure 2.8. Specialized roles of sensory neurons.....	36
Figure 2.9. Interneuron network for learning.....	38
Figure 2.10. Dopamine	40
Figure 2.11. Molecular mechanisms of learning	43
Figure 2.12. Protein synthesis.....	44
Figure 2.13. Schematic of network for learning	46
Figure 3.1. <i>C. elegans</i> displays transition in search behavior.....	68
Figure 3.2. Abrupt transitions	71
Figure 3.3. Animal behavior matches Infotaxis.....	74
Figure 3.4. Infotaxis is reducible to a drift-diffusion model.....	77
Figure 3.5. Proposed neural model	78
Table 2.1. Strain list.....	57
Table 2.2. Search numbers.....	59

ACKNOWLEDGMENTS

I would like to acknowledge Professor Tatyana Sharpee as well as Professor Sreekanth Chalasani for their mentorship and guidance. Dr. Sharpee transformed me from a mathematician concerned with ephemeral structures into a scientist, concerned with the physical world. Dr. Chalasani somehow managed to transform me into a biologist.

I would also like to acknowledge all the fantastic fellow graduate students and postdocs who I have shared a lab with, who I have learned more from than anyone else in my life.

My committee has also been helpful in their guidance and suggestions, Drs. William Kristan, Jr., Jing Wang, Charles Stevens, and Terrence Sejnowski.

Perhaps most important are the fellow students and friends who have kept me sane, and the graduate program for being full of such vitality and warmth.

Chapter 2, in full, is a manuscript in revision. Adam J Calhoun, Navin Pokala, James AJ Fitzpatrick, Tatyana O Sharpee, Sreekanth H Chalasani. Circuit mechanisms for evaluating uncertainty in *Caenorhabditis elegans*. In revision, 2014. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is a manuscript in preparation for submission in 2014. Adam J Calhoun, Sreekanth H Chalasani, Tatyana O Sharpee. Emergent decision-making in information maximizing behavior. In preparation, 2014. The dissertation author was the primary investigator and author of this paper.

VITA

2007 Master of Mathematics, First Class with Honours, University of St. Andrews; St. Andrews, Kingdom of Fife, Scotland

2014 Doctor of Philosophy, University of California, San Diego; La Jolla, CA

PUBLICATIONS

Calhoun AJ, Chalasani SH, Sharpee TO. Emergent decision-making in information maximizing behavior. *In preparation*.

Calhoun AJ, Pokala N, Fitzpatrick JAJ, Sharpee TO, Chalasani SH. Circuit mechanisms for evaluating uncertainty in *Caenorhabditis elegans*. *In revision*.

Sharpee TO, **Calhoun AJ**, Chalasani SH. Information theory of adaptation in neurons, behavior and mood, *Current Opinion Neurobiology*, 25, 47-53 (2014).

Reinagel P, Mankin E, **Calhoun AJ**. Speed and accuracy in a visual motion discrimination task as performed by rats, *arXiv*, arXiv:1206.0311.

Nandy D, **Calhoun A**, Windschitl J, Canfield RC, Linton MG. A new technique for measuring the twist of photospheric active regions without recourse to the force-free-field equation: reconfirming the hemispheric helicity trend, *Bulletin of the American Astronomical Society*, Vol 39, 128 (2007).

FIELDS OF STUDY

Major Field: Neuroscience (specialization in Computational Neuroscience)

ABSTRACT OF THE DISSERTATION

Mechanisms for optimal decision making in small neural circuits

by

Adam J Calhoun

Doctor of Philosophy in Neurosciences (specialization in Computational Neuroscience)

University of California, San Diego, 2014

Professor Tatyana Sharpee, Chair
Professor Terrence Sejnowski, Co-Chair

The need to acquire information about the variability in the world is paramount to optimal behavior yet it is not understood how this occurs on long timescales. Using the

nematode *Caenorhabditis elegans*, I developed a novel learning paradigm in which the animals utilize their previous experience to guide their exploration for new sources of food. Using a dimensionality reduction technique, I find that the animals are searching based on the variability in observed food, and the time course over which they are learning. I also investigate the neural circuitry that underlies this behavior and the mechanisms by which plasticity occurs – via dopamine and CREB. Dopamine acts on two distinct D1-like dopamine receptors, one on a sensory neuron and the other on its postsynaptic interneuron where CREB is also acting. The amount of CREB in the cell controls the rate at which the animal learns about the environment. Further, the sensory neurons which detect this variability are specialized for the task and only respond to large fluctuations in observed bacteria.

It is additionally unclear how to optimally use the available information. I utilize a model of optimal information-seeking behavior to show that the optimal behavior seeks reward in the local area for a finite amount of time before moving to a new area when the reliability of information about the current environment becomes low. Despite the model using a local decision rule, there is an emergent global change in behavior from local to global search. This optimal behavior can be approximated by a drift-diffusion model of decision-making, suggesting a deep connection between previous models of optimality and psychological theories of decision-making. Crucially, the behavior predicted by the model matches the behavior observed in *C. elegans*.

1. Introduction

All organisms must extract information from the environment and transform it into a useful signal to be able to make optimal decisions. *How* an organism extracts complex information about the environment is not well understood, especially when it is encoded in experiences that happen over a long time period. This dissertation attempts to understand one aspect of how a nervous system extracts information and learns about the environment, and how it optimally *uses* that information. I begin by exploring how information is acquired by the nervous system and follow with the biology of the organism under study. I then explain the theoretical framework to describe optimal behavior.

1.1 Circuits to extract information

Some of the circuit mechanisms that extract information are known; I begin by examining the visual system, which is perhaps the most well-studied sensory systems in neuroscience. Information processing begins in the retinal ganglion cells, which extract patterns of light and dark before sending that information further into the neural system (G. D. Field et al., 2010). As information flows through the nervous system, representations of the visual world become progressively more complex and the activity of cells systematically more decorrelated (Atick & Redlich, 1992). In area V1 of visual cortex, this information is transformed into representations of edges; interestingly, this representation is a function of the statistics of the natural world (Bell & Sejnowski, 1997; Olshausen & Field, 1997). Further in, V2 is believed to respond to

visual textures (Freeman, Ziemba, Heeger, Simoncelli, & Movshon, 2013) and V4 to translationally invariant edges (Nandy, Sharpee, Reynolds, & Mitchell, 2013), until reaching regions that represent faces or more complex objects (Tsao, Schweers, Moeller, & Freiwald, 2008). Yet how is this information transformed from the pointillist dots that photoreceptors represent?

At the level of the retina, it is clear that the precise pattern of connectivity from individual photoreceptors spatially cluster to produce the properties seen in individual ganglion cells (G. D. Field et al., 2010). The more complex properties of simple cells in V1 are similarly thought to arise from the arrangement of spatially aligned retinal ganglion cells pooling their connectivity to produce edges (HUBEL & WIESEL, 1962). Work here and in other sensory systems illustrates the importance of understanding the features that individual cells are extracting about the natural environment, and about the precise connectivity that transforms these features.

The nervous system does not just extract the presence or absence of dots and edges but also higher-order statistics such as luminance and contrast. At a short temporal timescales, these statistics are relatively straightforward to extract. For instance, the magnitude of stimulus intensity can be encoded at the level of photoreceptors or olfactory receptors (Choi et al., 2005). Yet the retina adapts to many orders of statistics in the visual scenery (Smirnakis, Berry, Warland, Bialek, & Meister, 1997) and visual cells adapt to all orders of timescales (Lundstrom, Higgs, Spain, & Fairhall, 2008). The retina is able to adapt its responses to different levels of contrast (Baccus & Meister, 2002), which it manages to achieve through a variety of mechanisms from synaptic changes (Jarsky et al., 2011) changes in internal integration

in each cell (Garvert & Gollisch, 2013). Visual cells do not just detect variance in light intensity, but also in other modalities such as motion and do so in an optimal manner (Fairhall, Lewen, Bialek, & de Ruyter Van Steveninck, 2001). How might these occur? One possibility might be via ion channel conductance changes from synaptic input (Chance, Abbott, & Reyes, 2002).

Yet these mechanisms all happen on relatively short timescales. Extracting information about the environment over longer timescales poses a much more challenging problem and may more simply be referred to as *learning*. One system that has been well-studied to extract some average value over these timescales is the reward-learning function of the basal ganglia. Here, value has been hypothesized to be encoded through the release of dopamine, particularly from the ventral tegmental area (VTA). It is commonly thought that dopamine released from the VTA represents a ‘rewarding’ signal (Wise, 2004) though it also has a role in motor control. Dopamine is released onto two classes of dopamine receptors, D1-like and D2-like, which modify cellular levels of cyclic AMP and are classically believed to represent potentiating or depressing modulatory pathways, respectively (Stoof & Keibadian, 1981). However this metaphor is complicated as D2-like receptors often function as autoreceptors to control the amount of dopamine released from presynaptic terminals (Hahn, Kullmann, Horn, & Levitan, 2006). Again, it is clear that a precise understanding of not just receptor type, but receptor action is required to fully understand these systems.

In order to encode average value, VTA releases dopamine in response to unexpected, but not expected, rewards or stimuli that predict reward (Schultz &

Hollerman, 1998). This signal activates D1-like dopamine receptors that in turn activate protein kinase A, phosphorylating the transcription factor CREB (Dudman et al., 2003; Kandel, 2001). Together, this suggests that dopamine is not representing reward itself, but rather predictions of reward. In fact, this matches the theoretical predictions of the Rescorla-Wagner temporal difference learning model (Sutton & Barto, 1990) which has been influential in guiding interpretations of dopamine function. Numerous models have suggested that a variety of learned behaviors, such as foraging preference, can arise from this simple algorithm (Montague, Dayan, Person, & Sejnowski, 1995; Niv, Joel, Meilijson, & Ruppin, 2002). Indeed, such algorithms have been wildly successful in robotics to learn what the appropriate behavior is for a given environment. I note that the role of dopamine in learning does not lie solely as an error signal, however. While the phasic response of neurons is implicated as a potential learning signal, dopamine neurons also display a tonic response to stimuli that do not represent error signals. Rather, it may represent response vigor (Niv, Daw, Joel, & Dayan, 2007; Schultz & Hollerman, 1998) or the reliance on learned information over exploration of new information (Beeler, 2012). Although this neural system is specifically believed to learn average value across time, it is interesting to note that there do exist neurons in the septum that respond to value uncertainty (Monosov & Hikosaka, 2013). However, it is not yet known what circuitry extracts this information from the environment and whether it is the same that learns mean value. In Chapter 2, I show a dopamine circuit that learns variability, a possible motif that may be observed in some form in the basal ganglia.

Despite all this, the number of neurons in the mammalian brain make understanding the complete neural circuits that extract information and produce behavior difficult to fully dissect. The most successful approaches to understanding these full neural networks come from invertebrates. One of the most successful comes from the work of Eric Kandel on the sea snail *Aplysia*, an animal with a much smaller nervous system than seen in mammals (~20,000 neurons) (Kandel, 2001). By focusing on the adaptation of gill withdrawal to mechanical stimulation, he found some of the first neural mechanisms that displayed learning. This circuit releases serotonin in response to mechanical stimulation, which in turn modulates the transmitter release from the gill sensory neuron (Kandel, 2001).

Other circuits have been useful at understanding the structure of neuronal circuits, chief among these being the lobster stomatogastric ganglion (STG). This circuit controls the movement of the crustacean gut through a series of rhythmic movements that grinds the food. Initially described by (Maynard & Dando, 1974), the STG has several intriguing features. First, it consists of a defined set of neurons whose activity continues when physically removed from the animal (Marder & Bucher, 2007). Second, the activity is controlled by sensory input that releases a large number of neuromodulators and neuropeptides. Finally, it is possible to record from most neurons in the circuit simultaneously.

One would think that with this type of model system, it should be possible to understand how the whole of a neural network functions. Despite this, the circuit remains something of a mystery. Several powerful computational studies have shown that highly disparate circuit parameters can yield functionally similar neuronal activity

which makes understanding structure alone insufficient to understand activity patterns (Prinz, Bucher, & Marder, 2004). Yet the output of these neuronal circuits is relatively stable suggesting that other mechanisms may play an important role in guiding network function. One key mechanism that has been identified is network modulation. Early on, (Turrigiano, Abbott, & Marder, 1994) identified activity-dependent homeostatic mechanisms as important to maintaining internal balance with subsequent work (O'Leary, Williams, Caplan, & Marder, 2013; Schulz, Goaillard, & Marder, 2006) showing that these mechanisms induce gene expression patterns that are highly correlated. This suggests that while the networks are highly plastic in their response to many neuromodulators and neuropeptides, there are in fact considerable constraints on the possible networks.

Although these are some of the most complete circuits, behavioral circuits have been identified and dissected in other systems. For example, eye-blink conditioning in the rodent brainstem (Nelson, Krispel, Sekirnjak, & Lac, 2003; Sekirnjak, Vissel, Bollinger, Faulstich, & Lac, 2003), escape circuitry in fish (Korn & Faber, 2005), and swimming in leeches (Briggman & Kristan, 2008). Imaging in leech during fictive swimming and crawling has revealed populations of neurons contributing to different stages of each behavior (Briggman & Abarbanel, 2005). Interestingly, these are not distinct networks for each behavior but rather overlapping networks (Briggman & Kristan, 2006), something seen in other invertebrate networks (Biron, Wasserman, Thomas, Samuel, & Sengupta, 2008).

1.2 Biology of *C. elegans*

In order to study the neural circuits that implement learning of these types of features, I have turned to the organism *Caenorhabditis elegans*. *C. elegans* is a soil-dwelling nematode that is primarily found among rotting fruit and compost (Brenner, 1974). Since it was described by (Brenner, 1974), it has become a model organism for a wide-range of biological applications. Among other advances, the complete neuronal circuit of the organism has been described; all 302 neurons in the hermaphrodite have systematically been identified, named, and their connections mapped (Varshney, Chen, Paniagua, Hall, & Chklovskii, 2011; White, Southgate, Thomson, & Brenner, 1986). This has allowed the organism to become a tractable starting point for investigations into minimal neuronal circuits for behavior.

Probably the most well-studied behavior in *C. elegans* is chemotaxis toward a chemical. This strategy combines principles of random walks with constantly updated information. At its most basic, this strategy is implemented in bacteria such as *E. coli* where it was first studied. Termed the run-and-tumble, bacteria will switch the flagella powering their movement from counter-clockwise to clockwise movement, switching them from straight swimming to tumbling (reorientation) (Berg, 2004). This allows the organism to continually reassess the environment despite loss of directional memory due to diffusion. *C. elegans* displays a similar – though slightly more complex – strategy. When chemotaxing toward an attractive substance, animals suppress their reorientations while going up a chemical gradient and increase their reorientations when going down a chemical gradient (Pierce-Shimomura, Morse, & Lockery, 1999), a chemotactic behavior is guided by seven known chemosensory

neurons (Bargmann, 2006). This is not the only search mechanism available to it: it can also use the ‘weathervane’ approach (Iino & Yoshida, 2009), wherein it will bias its movement slightly in the direction that the maximal chemical gradient can be found. These strategies are not mutually exclusive. Indeed, the weathervane approach is typically found when there are low changes in gradient while reorientations (run and tumble) occur in high gradient environments (Iino & Yoshida, 2009). This behavior is driven by sensory neurons that respond to changes in odorant concentration (Thiele, Faumont, & Lockery, 2009), allowing the animal to head toward or away from that odor depending on its attractiveness. In addition to chemotaxis, other behaviors such as *thermotaxis* operate similarly: increasing turns when going away from the optimal temperature, decreasing turns when going toward it (Bargmann & Mori, 1997). The traditional pathway for thermotaxis is through the sensory neuron AFD, though the olfactory sensory neuron AWC is also known to contribute (Biron et al., 2008). The temperature that animals thermotax toward is not fixed, but is dependent on the animal’s experience. This memory is not encoded deep in a neural circuit but rather directly in the sensory neuron AFD (Nishida, Sugi, Nonomura, & Mori, 2011).

Interestingly, the attractiveness or repulsiveness of an odor is a function of the precise neuron that detects it. After (Sengupta, Chou, & Bargmann, 1996) cloned the *odr-10* receptor that detects diacetyl, (Troemel, Kimmel, & Bargmann, 1997) expressed it on a neuron associated with aversive behavior, ASH. In so doing, the odor went from being attractive to being repulsive, indicating that chemotaxis is driven by the specific neurons activated by the particular chemical.

So far, the full extent of behavioral networks in *C. elegans* have been slow to appear. Initial investigations focused on the neuronal networks that controlled chemotaxis (Bargmann & Horvitz, 1991). These are characterized by initial chemosensory neurons (as described previously) that synapse primarily onto the interneurons AIA, AIY, AIB, and AIZ (Bargmann & Horvitz, 1991; Bargmann, Hartwig, & Horvitz, 1993; Tsalik & Hobert, 2003).

Beyond this, several other behavioral circuits have been mapped. The most complete diagram was investigated in (Gray, Hill, & Bargmann, 2005). Although it substantially overlapped with the chemotaxis circuit, it utilized a distinct subset of three primary sensory neurons (AWC, ASI, ASK). Further refinement of this circuit by (Chalasanani et al., 2007) investigated the neuronal dynamics as well as the synaptic sign of some of the neuronal connections. The first circuit for learning was described by (Ha et al., 2010). This circuit controlled the ability to learn about aversive (pathogenic) bacteria. Notably, it also contained many of the same interneurons found in other circuits – AIY, AIZ, and AIB. Similarly, a circuit mediating aversion to the chemical 1-octanol also contains these interneurons though it is extensively modified by aminergic and peptidergic modulation (Komuniecki, Harris, Hapiak, Wragg, & Bamber, 2011). The remarkable centrality of these neurons suggests that they play a key role in all locomotory behavior. However, other circuits find different neurons to be of more importance. In (Noble, Stieglitz, & Srinivasan, 2013), for instance, a simple four-neuron circuit consisting of two sensory neurons (AWB and ADF) and two interneurons (URX and RIC) described neuronal control of body fat through serotonin and octopamine.

The central behavior studied in this dissertation involves the search for new food. Many animals utilize a strategy known as area restricted search (ARS) to find new food by searching near where they have seen food in the past (Hills, Kalff, & Wiener, 2013; Sommerfeld, Kato, Ropert-Coudert, Garthe, & Hindell, 2013). *C. elegans* performs ARS via emission of many reorientation events immediately upon removal from food, a strategy that is controlled by dopamine and glutamate (Hills, Brockie, & Maricq, 2004). While the neural circuit describing this behavior is known (Gray et al., 2005; Wakabayashi, Kitagawa, & Shingai, 2004), it was previously not known if the behavior is modulated by prior experience.

Chapter 2 of this dissertation explores a circuit that modifies behavior based on previous experience of food distribution. This circuit contains a subset of neurons found in (Gray et al., 2005) as well as additional (dopaminergic) neurons. It is interesting to note, however, that certain chemosensory neurons that were previously found to be primary causes of behavior (Chalasanani et al., 2007) were not contributing to behavioral plasticity. This suggests that subcircuits may be fine-tuned to detect informative signals to modify behavior.

1.3 Optimal behavior

Once a neural network has acquired information, it must choose an appropriate action. What is the appropriate action to take given some set of knowledge? An influential model for optimal behavior utilizes the framework of Bayesian inference. This scheme assumes that an individual assigns probabilities to each event, representing the relative *belief* in that event. Then any new information can be

integrated into the organism's beliefs through Bayes' rule:

$$p(y|x) = \frac{p(x|y) p(y)}{p(x)}$$

There is extensive evidence that perception and action are consistent with Bayesian optimality. When humans attempt to infer, say, the speed of an object, they have a prior belief about how the world is structured (Tenenbaum & Griffiths, 2001). In this case, since we live in a world with more slowly-moving objects than quickly-moving objects, when forced to make a guess about a noisy object we tend to see such objects as moving more slowly than they are (Stocker & Simoncelli, 2006). Other work suggests that people's prior beliefs about the world combine in a Bayesian-optimal manner for predictions as diverse as movie run times and grosses, poem lengths, and cake baking times (Griffiths & Tenenbaum, 2006).

Yet representing probabilities seems like it may be difficult for a neural system. (Ma, Beck, Latham, & Pouget, 2006) proposed that the Poisson-like variability seen in cortex makes these sorts of computations tractable. But the large number of neurons required seems difficult for organisms with smaller numbers of neurons, despite the idea that they, too, act as Bayesian-optimal agents. In Chapter 3, I show a Bayesian-optimal behavior that can be approximated by a one-dimensional model.

A related optimality criterion is information-maximization. This hypothesis states that the nervous system attempts to maximize its information about the world (Rieke, Warland, & De Ruyter van Steveninck, 1997). The mutual information between a neuronal response and the stimulating environment can be computed as the

difference between the entropy of the neural response $H(r)$ and the average entropy $\langle H(r|stim) \rangle$. In other words:

$$I(r; stim) = H(r) - \langle H(r|stim) \rangle_{stim}$$

Indeed, sensory neurons appear to maximize their information about the world (Barlow, 1961; Fairhall et al., 2001), as do networks (Tkacik, Prentice, Balasubramanian, & Schneidman, 2010). To extend this analogy, it is possible that if sensory neurons are maximizing their information about stimuli, perhaps organisms are maximizing their information about reward.

This is the basis for the formulation of chemotaxis-like behavior termed ‘infotaxis’ (Najemnik & Geisler, 2005; Vergassola, Villermaux, & Shraiman, 2007).

The goal is to maximize the function:

$$\begin{aligned} \Delta S(r_{new}|r_{current}) \\ = -P(r_{current})S_{current} + (1 - P(r_{current})) \sum_{n=0}^a p_n(r_{new})\Delta S_n(r_{new}) \end{aligned}$$

Here, $P(r)$ describes the probability of finding an odor source at location r , the entropy of the distribution is denoted as S and the current position $r_{current}$. The terms p_n represent the probability of n odorant detections where ΔS_n is the expected change in entropy from gaining those detections. This is distinct from a pure chemotaxis strategy where the searcher attempts to maximize the number of detection events $\sum_{n=0}^a np_n$.

With this strategy, animals use all available information to guide their search behavior and detection of single particles are sufficient to accurately guide behavior. As the odorant level increases, the large amount of incoming information converges the behavior to chemotaxis making these strategies equivalent at high odor levels.

Interestingly, this strategy explains the characteristic properties of moth flight, such as cross-wind zigzagging far from pheromonal sources and gradient ascent near those sources. This approach is not unique to chemosensation, however, as it makes strong predictions about other behaviors such as eye movement (Najemnik & Geisler, 2005) and possibly even retinal responses to single photons (Schwartz & Rieke, 2013).

A final model of optimal behavior is found in the evidence-accumulation model. In this model, every new input provides information to some decision variable that represents the accumulated knowledge about some task: whether to continue, whether option A or option B is better, and so on (Ratcliff & McKoon, 2008). This model of behavior has explained a large number of cognitive findings in the field. And indeed, this is an optimal strategy when forced to make a time/accuracy tradeoff (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006).

Although initially proposed as a purely psychological description of behavior, this model has found support in neuroscience. Neural correlates of decision-variables have been found in multiple areas of the brain. The most well-studied example is in the lateral intraparietal cortex (area LIP) in response to randomly moving dots (Shadlen & Newsome, 1996). In this task, some percentage of dots move in a uniform direction while the rest move randomly. At each time step, the particular dot that moves coherently or randomly differs. Subjects are then forced to decide whether the dots are moving, say, left or right. The difficulty of this task varies by the proportion of dots that move coherently. When recording from LIP, neurons respond by slowly increasing (or decreasing) their firing rate in a noisy manner until reaching a 'threshold', upon which an action (decision) occurs (Shadlen & Newsome, 2001).

Crucially, the speed at which the firing rate increases is determined by the difficulty of the task, as the DDM would suggest.

However, it remains to be shown what the optimal evidence accumulation rate is for large classes of problems. In Chapter 3, we describe *C. elegans* behavior as consistent with an infotactic search approach. We then show that implicit in the model is a decision-variable that can be modeled as a DDM. This presents a powerful argument for the optimality of DDMs, and ties together the framework with Bayesian optimality.

1.4 References

- Atick, J. J., & Redlich, A. N. (1992). What Does the Retina Know about Natural Scenes? *Neural Computation*, 4(2), 196–210. doi:10.1016/0042-6989(72)90139-3
- Baccus, S. A., & Meister, M. (2002). Fast and slow contrast adaptation in retinal circuitry. *Neuron*, 36(5), 909–919.
- Bargmann, C. (2006). Chemosensation in *C. elegans*. *WormBook*. doi:10.1895/wormbook.1.123.1
- Bargmann, C. I., & Horvitz, H. R. (1991). Chemosensory neurons with overlapping functions direct chemotaxis to multiple chemicals in *C. elegans*. *Neuron*, 7(5), 729–742.
- Bargmann, C. I., & Mori, I. (1997). 25 Chemotaxis and Thermotaxis. *Cold Spring Harbor Monograph Archive*.
- Bargmann, C. I., Hartwig, E., & Horvitz, H. R. (1993). Odorant-selective genes and neurons mediate olfaction in *C. elegans*. *Cell*, 74(3), 515–527.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. *Sensory Communication*.
- Beeler, J. A. (2012). Thorndike's Law 2.0: Dopamine and the Regulation of Thrift. *Frontiers in Neuroscience*, 6, 116. doi:10.3389/fnins.2012.00116

- Bell, A. J., & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Research*, *37*(23), 3327–3338. doi:10.1016/S0042-6989(97)00121-1
- Berg, H. C. (2004). *E. coli in Motion* - Google Books.
- Biron, D., Wasserman, S., Thomas, J. H., Samuel, A. D. T., & Sengupta, P. (2008). An olfactory neuron responds stochastically to temperature and modulates *Caenorhabditis elegans* thermotactic behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(31), 11002–11007. doi:10.1073/pnas.0805004105
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. doi:10.1037/0033-295X.113.4.700
- Brenner, S. (1974). THE GENETICS OF CAENORHABDITIS ELEGANS.
- Briggman, K. L., & Kristan, W. B. (2006). Imaging dedicated and multifunctional neural circuits generating distinct behaviors. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *26*(42), 10925–10933. doi:10.1523/JNEUROSCI.3265-06.2006
- Briggman, K. L., & Kristan, W. B., Jr. (2008). Multifunctional Pattern-Generating Circuits. *Annual Review of Neuroscience*, *31*(1), 271–294. doi:10.1146/annurev.neuro.31.060407.125552
- Briggman, K., & Abarbanel, H. (2005). Optical Imaging of Neuronal Populations During Decision-Making. *Science (New York, NY)*.
- Chalasan, S. H., Chronis, N., Tsunozaki, M., Gray, J. M., Ramot, D., Goodman, M. B., & Bargmann, C. I. (2007). Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*. *Nature*, *450*(7166), 63–70. doi:10.1038/nature06292
- Chance, F. S., Abbott, L. F., & Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron*, *35*(4), 773–782.
- Choi, S.-Y., Borghuis, B. G., Borghuis, B., Rea, R., Levitan, E. S., Sterling, P., & Kramer, R. H. (2005). Encoding light intensity by the cone photoreceptor synapse. *Neuron*, *48*(4), 555–562. doi:10.1016/j.neuron.2005.09.011
- Dudman, J. T., Eaton, M. E., Rajadhyaksha, A., Macías, W., Taher, M., Barczak, A., et al. (2003). Dopamine D1 receptors mediate CREB phosphorylation via phosphorylation of the NMDA receptor at Ser897-NR1. *Journal of*

Neurochemistry, 87(4), 922–934.

- Fairhall, A. L., Lewen, G. D., Bialek, W., & de Ruyter Van Steveninck, R. R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature*, 412(6849), 787–792. doi:10.1038/35090500
- Field, G. D., Gauthier, J. L., Sher, A., Greschner, M., Machado, T. A., Jepson, L. H., et al. (2010). Functional connectivity in the retina at the resolution of photoreceptors. *Nature*, 467(7316), 673–677. doi:10.1038/nature09424
- Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (2013). A functional and perceptual signature of the second visual area in primates. *Nature Neuroscience*, 16(7), 974–981. doi:10.1038/nn.3402
- Garvert, M. M., & Gollisch, T. (2013). Local and global contrast adaptation in retinal ganglion cells. *Neuron*, 77(5), 915–928. doi:10.1016/j.neuron.2012.12.030
- Gray, J. M., Hill, J. J., & Bargmann, C. I. (2005). A circuit for navigation in *Caenorhabditis elegans*.
- Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, 17(9), 767–773. doi:10.1111/j.1467-9280.2006.01780.x
- Ha, H.-I., Hendricks, M., Shen, Y., Gabel, C. V., Fang-Yen, C., Qin, Y., et al. (2010). Functional organization of a neural network for aversive olfactory learning in *Caenorhabditis elegans*. *Neuron*, 68(6), 1173–1186. doi:10.1016/j.neuron.2010.11.025
- Hahn, J., Kullmann, P. H. M., Horn, J. P., & Levitan, E. S. (2006). D2 autoreceptors chronically enhance dopamine neuron pacemaker activity. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 26(19), 5240–5247. doi:10.1523/JNEUROSCI.4976-05.2006
- Hills, T. T., Kalff, C., & Wiener, J. M. (2013). Adaptive Lévy processes and area-restricted search in human foraging. *PloS One*, 8(4), e60488. doi:10.1371/journal.pone.0060488
- Hills, T., Brockie, P. J., & Maricq, A. V. (2004). Dopamine and glutamate control area-restricted search behavior in *Caenorhabditis elegans*. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 24(5), 1217–1225. doi:10.1523/JNEUROSCI.1569-03.2004
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160,

106–154.

- Iino, Y., & Yoshida, K. (2009). Parallel use of two behavioral mechanisms for chemotaxis in *Caenorhabditis elegans*. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *29*(17), 5370–5380. doi:10.1523/JNEUROSCI.3633-08.2009
- Jarsky, T., Cembrowski, M., Logan, S. M., Kath, W. L., Riecke, H., Demb, J. B., & Singer, J. H. (2011). A synaptic mechanism for retinal adaptation to luminance and contrast. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *31*(30), 11003–11015. doi:10.1523/JNEUROSCI.2631-11.2011
- Kandel, E. R. (2001). The molecular biology of memory storage: a dialogue between genes and synapses. *Science (New York, NY)*, *294*(5544), 1030–1038. doi:10.1126/science.1067020
- Komuniecki, R., Harris, G., Hapiak, V., Wragg, R., & Bamber, B. (2011). Monoamines activate neuropeptide signaling cascades to modulate nociception in *C. elegans*: a useful model for the modulation of chronic pain? *Invertebrate Neuroscience : in.* doi:10.1007/s10158-011-0127-0
- Korn, H., & Faber, D. S. (2005). The Mauthner cell half a century later: a neurobiological model for decision-making? *Neuron*, *47*(1), 13–28. doi:10.1016/j.neuron.2005.05.019
- Lundstrom, B. N., Higgs, M. H., Spain, W. J., & Fairhall, A. L. (2008). Fractional differentiation by neocortical pyramidal neurons. *Nature Neuroscience*, *11*(11), 1335–1342. doi:10.1038/nn.2212
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–1438. doi:10.1038/nn1790
- Marder, E., & Bucher, D. (2007). Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs. *Annual Review of Physiology*, *69*, 291–316. doi:10.1146/annurev.physiol.69.031905.161516
- Maynard, D. M., & Dando, M. R. (1974). The structure of the stomatogastric neuromuscular system in *Callinectes sapidus*, *Homarus americanus* and *Panulirus argus* (Decapoda Crustacea). *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *268*(892), 161–220.
- Monosov, I. E., & Hikosaka, O. (2013). Selective and graded coding of reward uncertainty by neurons in the primate anterodorsal septal region. *Nature Neuroscience*, *16*(6), 756–762. doi:10.1038/nn.3398

- Montague, P. R., Dayan, P., Person, C., & Sejnowski, T. J. (1995). Bee foraging in uncertain environments using predictive hebbian learning. *Nature*, *377*(6551), 725–728. doi:10.1038/377725a0
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, *434*(7031), 387–391. doi:10.1038/nature03390
- Nandy, A. S., Sharpee, T. O., Reynolds, J. H., & Mitchell, J. F. (2013). The fine structure of shape tuning in area V4. *Neuron*, *78*(6), 1102–1115. doi:10.1016/j.neuron.2013.04.016
- Nelson, A. B., Krispel, C. M., Sekirnjak, C., & Lac, du, S. (2003). Long-lasting increases in intrinsic excitability triggered by inhibition. *Neuron*, *40*(3), 609–620.
- Nishida, Y., Sugi, T., Nonomura, M., & Mori, I. (2011). Identification of the AFD neuron as the site of action of the CREB protein in *Caenorhabditis elegans* thermotaxis. *EMBO Reports*, *12*(8), 855–862. doi:10.1038/embor.2011.120
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*, *191*(3), 507–520. doi:10.1007/s00213-006-0502-4
- Niv, Y., Joel, D., Meilijson, I., & Ruppin, E. (2002). Evolution of Reinforcement Learning in Uncertain Environments: A Simple Explanation for Complex Foraging Behaviors. *Adaptive Behavior*.
- Noble, T., Stieglitz, J., & Srinivasan, S. (2013). An integrated serotonin and octopamine neuronal circuit directs the release of an endocrine signal to control *C. elegans* body fat. *Cell Metabolism*, *18*(5), 672–684. doi:10.1016/j.cmet.2013.09.007
- O'Leary, T., Williams, A. H., Caplan, J. S., & Marder, E. (2013). Correlations in ion channel expression emerge from homeostatic tuning rules. *Proceedings of the ...*
- Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, *37*(23), 3311–3325. doi:10.1016/S0042-6989(97)00169-7
- Pierce-Shimomura, J. T., Morse, T. M., & Lockery, S. R. (1999). The fundamental role of pirouettes in *Caenorhabditis elegans* chemotaxis. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *19*(21), 9557–9569.
- Prinz, A. A., Bucher, D., & Marder, E. (2004). Similar network activity from disparate

- circuit parameters. *Nature Neuroscience*, 7(12), 1345–1352. doi:10.1038/nn1352
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922. doi:10.1162/neco.2008.12-06-420
- Rieke, F., Warland, D., & De Ruyter van Steveninck, R. (1997). *Spikes* | The MIT Press.
- Schultz, W., & Hollerman, J. R. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, 1(4), 304–309. doi:10.1038/1124
- Schulz, D. J., Goillaud, J.-M., & Marder, E. (2006). Variable channel expression in identified single and electrically coupled neurons in different animals. *Nature Neuroscience*, 9(3), 356–362. doi:10.1038/nn1639
- Schwartz, G. W., & Rieke, F. (2013). Controlling gain one photon at a time. *eLife*, 2, e00467. doi:10.7554/eLife.00467
- Sekirnjak, C., Vissel, B., Bollinger, J., Faulstich, M., & Lac, du, S. (2003). Purkinje cell synapses target physiologically unique brainstem neurons. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 23(15), 6392–6398.
- Sengupta, P., Chou, J. H., & Bargmann, C. I. (1996). odr-10 encodes a seven transmembrane domain olfactory receptor required for responses to the odorant diacetyl. *Cell*, 84(6), 899–909.
- Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National ...*
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86(4), 1916–1936.
- Smirnakis, S. M., Berry, M. J., Warland, D. K., Bialek, W., & Meister, M. (1997). Adaptation of retinal processing to image contrast and spatial scale. *Nature*, 386(6620), 69–73. doi:10.1038/386069a0
- Sommerfeld, J., Kato, A., Ropert-Coudert, Y., Garthe, S., & Hindell, M. A. (2013). Foraging parameters influencing the detection and interpretation of area-restricted search behaviour in marine predators: a case study with the masked booby. *PloS One*, 8(5), e63742. doi:10.1371/journal.pone.0063742

- Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*(4), 578–585. doi:10.1038/nn1669
- Stoof, J. C., & Keibian, J. W. (1981). Opposing roles for D-1 and D-2 dopamine receptors in efflux of cyclic AMP from rat neostriatum. *Nature*, *294*(5839), 366–368.
- Sutton, R. S., & Barto, A. G. (1990). *Time-derivative models of Pavlovian reinforcement*. The MIT Press.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *The Behavioral and Brain Sciences*, *24*(4), 629–40– discussion 652–791.
- Thiele, T. R., Faumont, S., & Lockery, S. R. (2009). The neural network for chemotaxis to tastants in *Caenorhabditis elegans* is specialized for temporal differentiation. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *29*(38), 11904–11911. doi:10.1523/JNEUROSCI.0594-09.2009
- Tkacik, G., Prentice, J. S., Balasubramanian, V., & Schneidman, E. (2010). Optimal population coding by noisy spiking neurons.
- Troemel, E. R., Kimmel, B. E., & Bargmann, C. I. (1997). Reprogramming chemotaxis responses: sensory neurons define olfactory preferences in *C. elegans*. *Cell*, *91*(2), 161–169.
- Tsalik, E. L., & Hobert, O. (2003). Functional mapping of neurons that control locomotory behavior in *Caenorhabditis elegans*. *Journal of Neurobiology*, *56*(2), 178–197. doi:10.1002/neu.10245
- Tsao, D. Y., Schweers, N., Moeller, S., & Freiwald, W. A. (2008). Patches of face-selective cortex in the macaque frontal lobe. *Nature Neuroscience*, *11*(8), 877–879. doi:10.1038/nn.2158
- Turrigiano, G., Abbott, L. F., & Marder, E. (1994). Activity-dependent changes in the intrinsic properties of cultured neurons. *Science (New York, NY)*, *264*(5161), 974–977.
- Varshney, L. R., Chen, B. L., Paniagua, E., Hall, D. H., & Chklovskii, D. B. (2011). Structural Properties of the *Caenorhabditis elegans* Neuronal Network. *PLoS Computational Biology*, *7*(2), e1001066. doi:10.1371/journal.pcbi.1001066
- Vergassola, M., Villermanx, E., & Shraiman, B. I. (2007). “Infotaxis” as a strategy for searching without gradients. *Nature*, *445*(7126), 406–409.

doi:10.1038/nature05464

Wakabayashi, T., Kitagawa, I., & Shingai, R. (2004). Neurons regulating the duration of forward locomotion in *Caenorhabditis elegans*. *Neuroscience Research*, *50*(1), 103–111. doi:10.1016/j.neures.2004.06.005

White, J. G., Southgate, E., Thomson, J. N., & Brenner, S. (1986). The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *314*(1165), 1–340. doi:10.1098/rstb.1986.0056

Wise, R. A. (2004). Dopamine, learning and motivation. *Nature Reviews Neuroscience*, *5*(6), 483–494. doi:10.1038/nrn1406

2. Neural mechanisms for evaluating uncertainty in *Caenorhabditis elegans*

2.1 Introduction

Animals and their underlying neural circuits respond to changing sensory environments by integrating prior experience with new information to drive appropriate behaviors. In particular, the consistency of the environment plays a crucial role in shaping an animal's behavioral strategy as unpredictable environments lead to unpredictable rewards. When presented with a choice, animals often prefer a behavioral strategy that will generate stable rewards to one with uncertain rewards (MacLean et al., 2012; Platt and Huettel, 2008). Perhaps selecting a strategy with uncertain reward may lead to an outcome without any reward; for example, in the case of food, no reward may represent possible starvation (Watson, 2008). Thus, animals must devote considerable resources to determining reliability in their environments (Escobar et al., 2007; MacLean et al., 2012). This task can be particularly challenging as the underlying neural circuit must evaluate the variance of reward in the environment over the timescale of minutes to hours and then guide a behavioral strategy that will itself last a similar timescale.

In the visual system, the rate and statistics of action potential firing have been shown to encode information about rapidly occurring variations in stimuli. The underlying cellular mechanism usually occurs within seconds. Moreover, these studies also show that speed of resolving ambiguities approaches the physical limits imposed by the sampling rate and noise (Fairhall et al., 2001; Wark et al., 2009). However, little is known about how a neural circuit evaluates reliability over many minutes and generates complex behaviors that last a similar timescale in response to a changing

environment . One unconventional method is to analyze these complex behaviors and the underlying minimal neural circuits in a simple, genetically tractable model. The relatively small *C. elegans* nervous system consisting of just 302 neurons with identified connections is ideally suited for analyzing these circuits at the resolution of single cells (Chalasani et al., 2007; Chalfie et al., 1985; de Bono and Maricq, 2005; White et al., 1986).

Many animals including *C. elegans* search by spending a large amount of time near a previously observed reward, a strategy termed area-restricted search (Gray et al., 2005; Hills et al., 2013; Sommerfeld et al., 2013; Thums et al., 2011; Wakabayashi et al., 2004). Animals removed from bacterial food (large reward) execute an initial ‘local search’ of a restricted area for about 15 minutes by interrupting forward movements by seemingly random reorientations (turns). This behavior is driven by three pairs of sensory neurons that respond to food: AWC, ASI and ASK. After 15 minutes, animals disperse by suppressing these reorientations, which is termed the ‘global search’. Moreover, cell ablation experiments have identified the entire 46-neuron circuit consisting of the above 6 sensory neurons, 14 interneurons and 26 motor neurons that regulate local search behavior (Gray et al., 2005; Wakabayashi et al., 2004). We show that this complex behavior involves each animal evaluating the distribution of food (variability of reward) and using that information to drive a complex local search lasting many minutes. We use a combination of genetics, behavioral analysis, imaging and theoretical methods to investigate how this neural circuit evaluates a bacterial patch and drives local search. We identify a novel circuit

motif that measures the variance in food and the underlying dopaminergic circuit that generates appropriate local search.

2.2 Results

2.2.1 Prior food experience modifies local search behavior

To test whether prior experience influenced search strategy we analyzed the behavior of animals exposed to different bacteria. We found that animals removed from lawns of three different bacteria, *Escherichia coli* strains OP50 and HB101 and *Bacillus megaterium* (DA1880) executed an initial local search and then transitioned to a global search state (Figure 2.1A). We observe a dramatic difference in the number of turns executed by animals in local search state, while those in global search mode are not influenced by prior food. Surprisingly, we find that animals removed from a low quality food source, *B. megaterium* (Avery and Shtonda, 2003), execute the most number of turns during local search (Figure 2.1B). These data suggest that prior food experience influences the number of turns executed by animals during local search.

To confirm the role of prior food experience we transferred animals from a lawn containing one bacteria to one containing the second bacteria for different time periods and analyzed the local search behavior. We hypothesized that if the animals were integrating information about their environment then they would switch their behavior to reflect the second bacteria over some period of time. Consistent with our hypothesis, we found that animals modified their local search to match those removed from the second bacteria after they explored a lawn of the second bacteria for 45

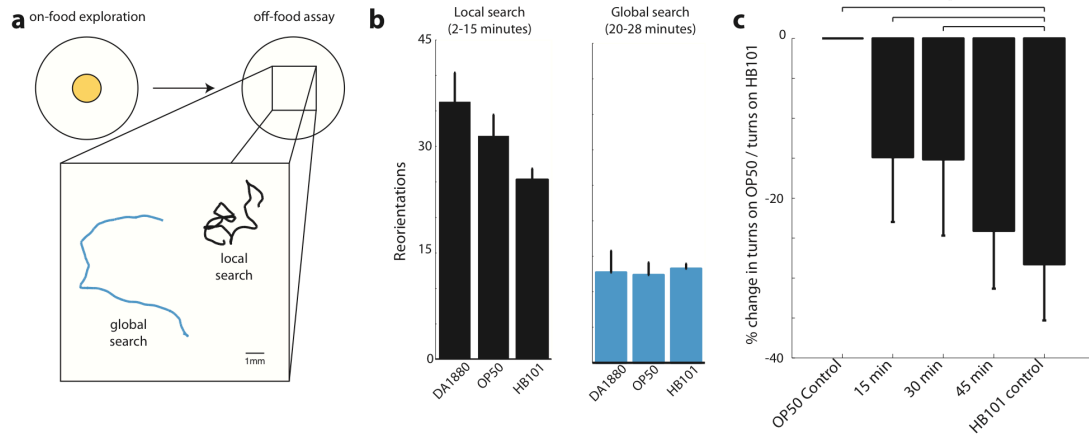


Figure 2.1. On food experience modifies off-food behavior. **a**, Representation of behavioral assay. Example track from a worm removed from food. **b**, The precise search strategy is dependent on the food that the worm has experienced, with some bacteria causing the worm to emit more turns over the first fifteen minutes and some to emit fewer. **c**, Worms transferred between different bacterial strains modify their behavior in under 45 minutes.

minutes (Figure 2.1C). These data indicate that animals use the information gained in the last 45 minutes of their prior food experience to modify their local search behavior.

To begin investigating whether the animals were simply responding to the type of food, we analyzed the behavior of animals removed from different sized lawns of the same bacteria. Surprisingly, we observe that the size of the bacteria patch that the animal was removed from is directly proportional to the number of reorientations that the animal executes during local search (Figure 2.2A, Figure 2.3). In particular, patch experience specifically modified the number of large-angled turns that cause large reorientations in locomotory paths, but not non-reorienting reversals (Figure 2.3). Moreover, the size of the bacteria lawn only influences local search behavior and has no influence on global search behavior (Figure 2.2B). This patch size learning behavior was not specific to *E. coli* patches, as animals removed from *Comamonas* sp. or *Pseudomonas fluorescens* also demonstrate a strong correlation between size of the bacteria patch and local search area confirming that *C. elegans* uses prior experience to modify local search as general strategy (Figure 2.2C).

2.2.2 Environmental variability can quantitatively predict local search

To gain insight into the sensory environment that the animal is experiencing, we measured the gradient of bacteria in both small and large patches by analyzing bacteria expressing green fluorescent protein (Labrousse et al., 2000). By using fluorescence intensity measurements as a proxy for bacterial concentration, we found

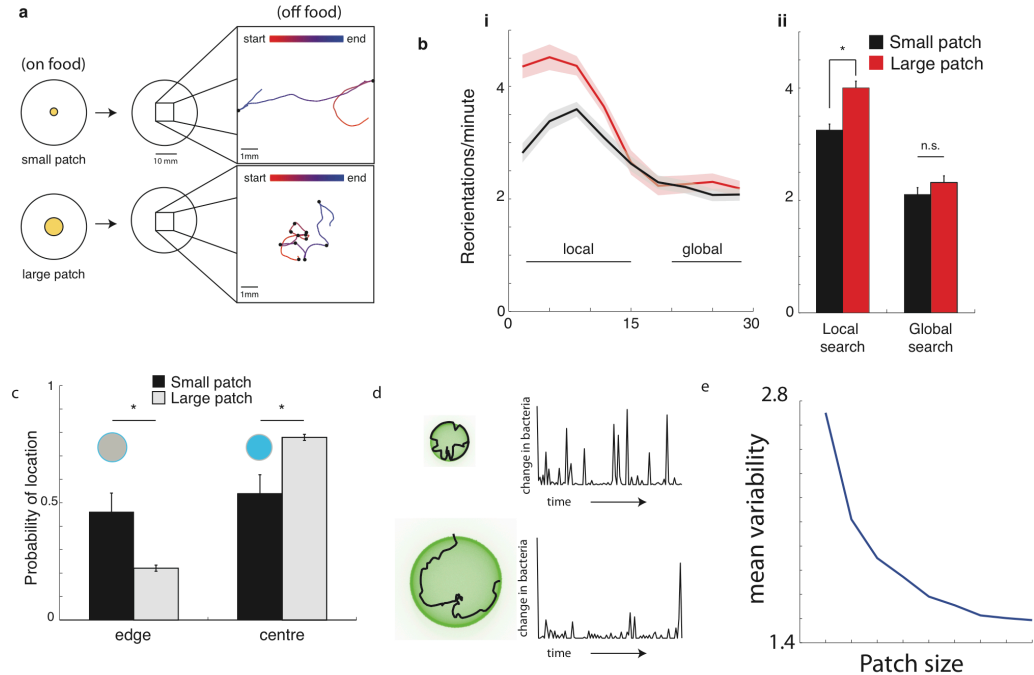


Figure 2.2. Worms on differently-sized food patches experience different sensory environments. a, Representations of small- and large-sized patches (left) and example tracks (right) of animals removed from small (0.7 cm in diameter) and large sized (1.9 cm in diameter) patches. Tracks show movement during first three minutes of local search (red at the start, blue at the end, black represent a reorientation), with the animal from the small patch performing two reorientations and a mean-squared diffusion (MSD) of 26 mm²/minute and from the large patch eleven turns and a MSD of 44.6 mm²/sec. b, (i,ii) Animals removed from a large patch perform more turns during local search (2-15 minutes) when compared to animals from a small patch (t-test, $p < 10^{-4}$). Global search (20-28 minutes) is not affected by patch size. c, Real worms on food show a similar decrease in residence time on the edge of the patch from small to large food patches. d, A markov model was run to simulate worm movement across patches. Movement probability between regions on the food patch was extracted from on-food tracking. e, In simulation, worms on smallpatches (black line) spend more time on the edge of the patch while worms on large patches spend less time. There is a marked decrease in the number of movements between edge and center of the patch from small to large (inset).

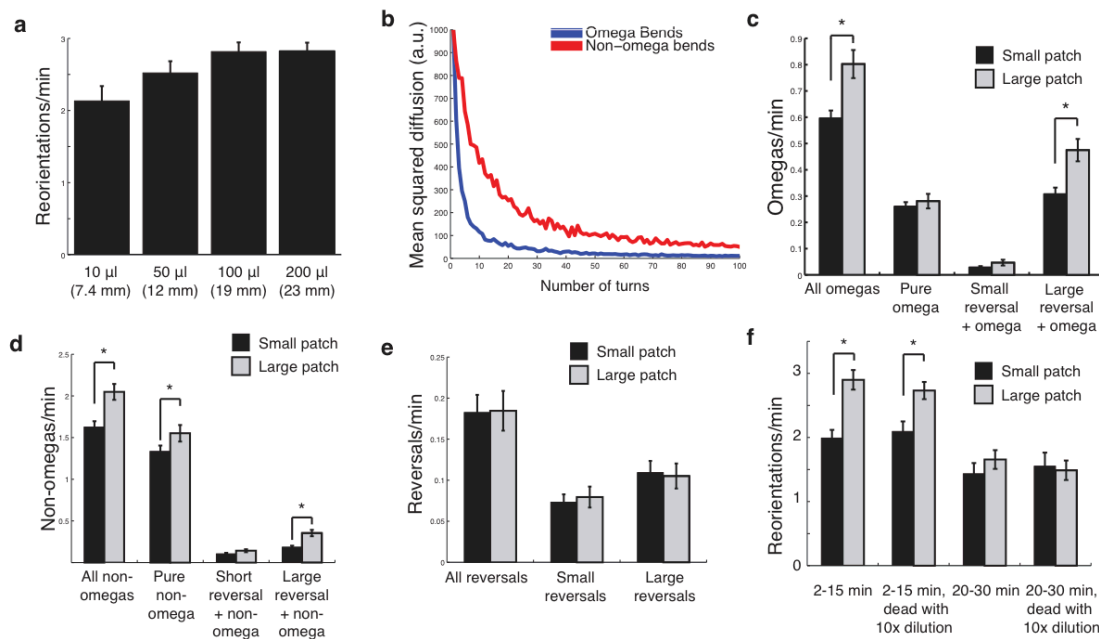


Figure 2.3. Specifics of learning. a, Increasing diameter of bacterial lawn leads to increasing numbers of turns upon removal (local search). b, the number of turns an animal executes is also inversely proportional to the area searched. c, Using simulations, we find that the numbers of omega and non-omega turns are inversely proportional to the area searched. d-f, Learning modifies omega and non-omega turns, but not reversals. g, Animals can discriminate between small and large patches of dead and diluted bacteria. h, Animals can also distinguish between different sized lawns of *Comamonas* and *Pseudomonas*. Error bars represent s.e.m. and * indicates significance (t-test, $p < 0.05$).

that patches of different sizes have a similar steep gradient at the edge and a flat centre (Figure 2.4). We next investigated how sensory experience may differ for animals exploring differently sized food patches by analyzing their behavior while on food. We find that animals spend more time at the edge of a small patch when compared to those on a large patch (Figure 2.2D). We then used a simple Markov model to simulate worm movement on patches of different sizes. Our simulation suggests that animals on a small patch are more likely to encounter the edge when compared to those exploring a large lawn (Figure 2.2E). Moreover, the probabilities predicted by the simulation matches those observed in real animals exploring food patches (Figure 2.4). These data suggest that animals exploring different food patches will experience different sensory environments.

Since animals encountering the edge of a food patch encounter more food, we examined whether animals taken from food patches where bacteria was diluted would behave differently. However, they appeared to use the same strategy (Figure 2.4). We hypothesized that animals exploring different sized patches may encounter different variability. We find that animals exploring a small patch experience greater variance in their environment compared to those on a large patch (Figure 2.2F). Further, the variability that a single worm experiences is predictive of its off-food behavior (Figure 2.5). Taken together these results indicate that animals exploring different sized lawns experience different variances in their environment.

In order to identify the best predictor of the number of turns executed during local search, we decoded the behavior of animals while exploring the patch. We applied a dimensionality reduction technique (Fitzgerald et al., 2011) to extract a

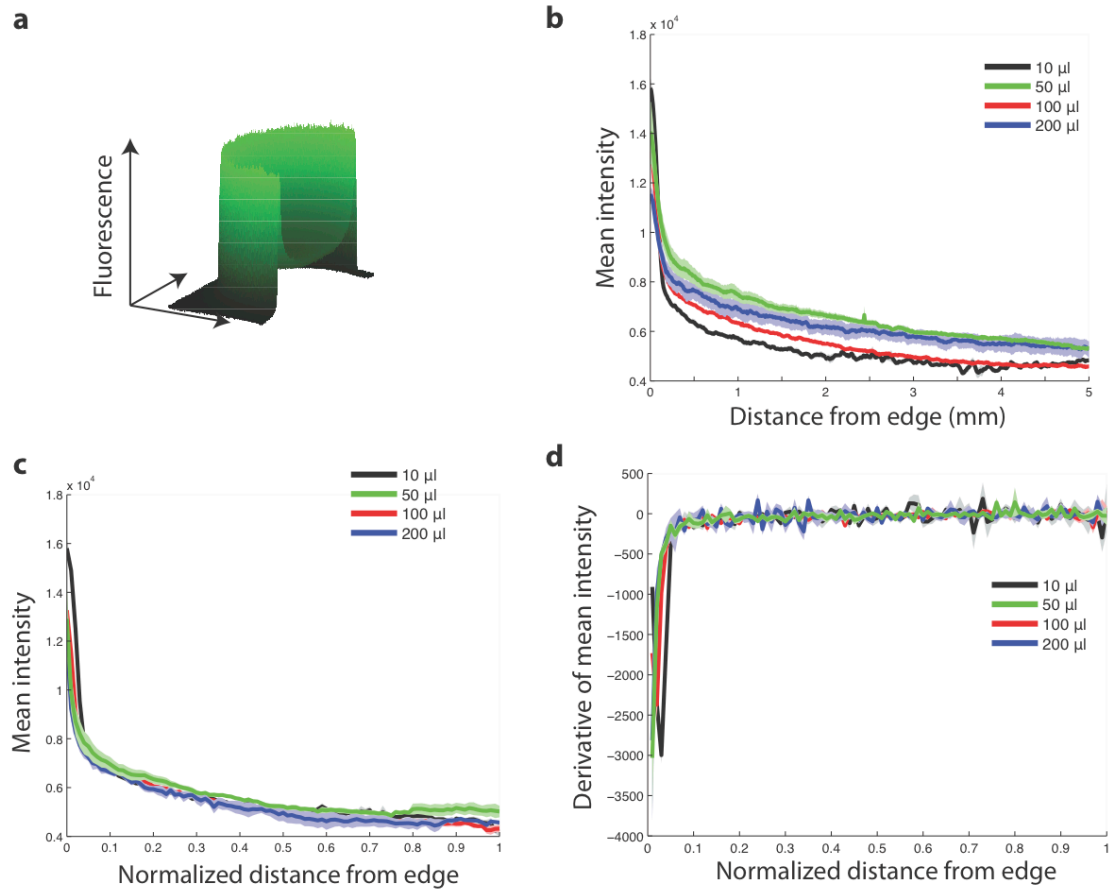


Figure 2.4. Measuring bacterial gradients. a, Cutaway of fluorescence values of small bacterial lawn expressing green fluorescent protein. b, Fluorescence levels are similar between large and small food patches on the edge. c, Fluorescence intensity normalized by distance from the edge is constant across bacterial patches of different sizes. d, The instantaneous gradient is high at the edge and roughly zero in the center. Shaded areas represent s.e.m.

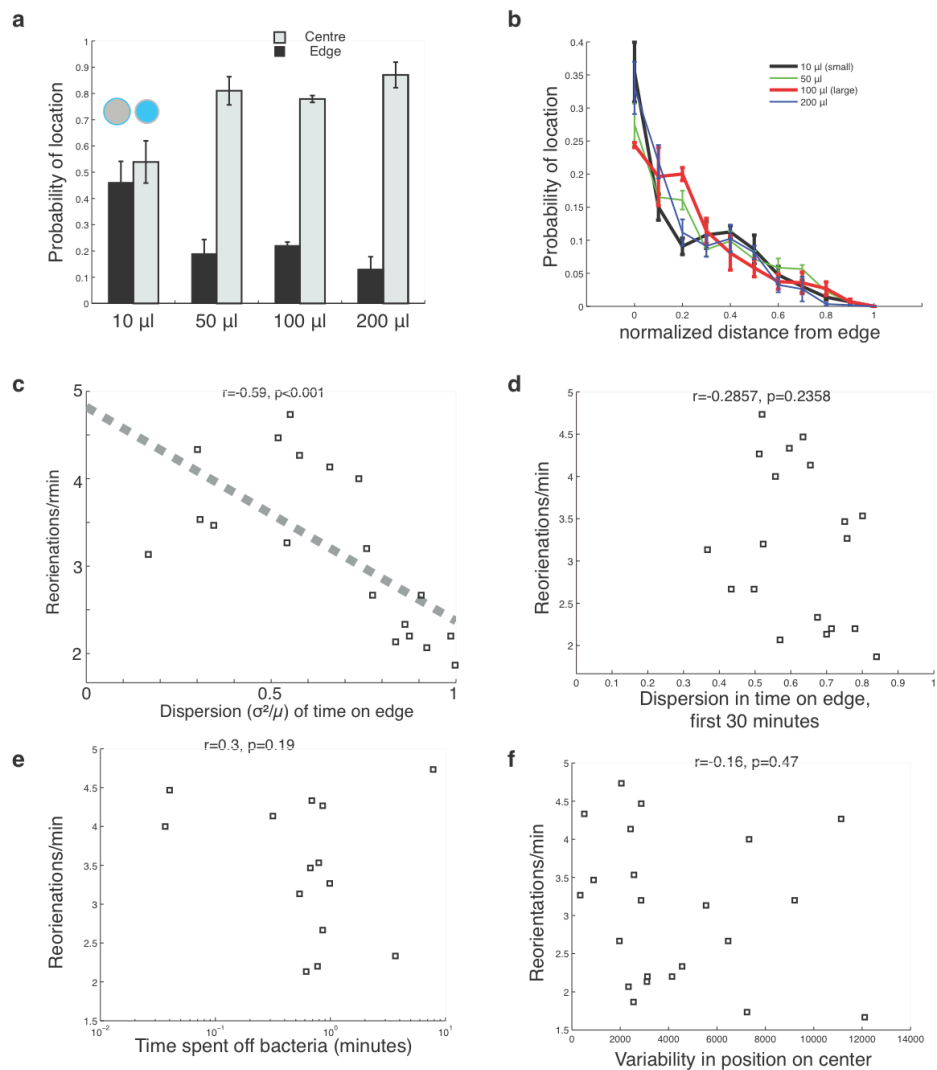


Figure 2.5. Quantifying on-food behavior. a, The amount of time spent on the edge decreases as patch size increases. b, When normalized by distance from the edge, animals spend an equivalent amount of time at each position regardless of the size of the bacterial patch. c, Dispersion of time on edge between thirty and sixty minutes prior to removal from food is not correlated with the number of turns that occur during local search. d, Neither time spent off the food patch nor e, variance of position across the center of the patch during thirty minutes prior to removal from food predict number of turns during local search. Error bars represent s.e.m.

behavioral filter. The filter has three notable aspects: the immediate five minutes preceding removal from food has no bearing on future behavior, nor does any time prior to thirty minutes before removal. However, the duration between 30 minutes and 5 minutes preceding removal accounts for a large degree of variability in the number of turns each animal will make after removal from food (Figure 2.6A). We also tested a number of other possible predictors of behavior, including time spent off food patch or movements on the center of a patch, and found that none of them could predict the turns during local search (Figure 2.5). More generally, we found that our behavioral filter could accurately predict the number of turns during local search using the variability experienced by an animal (Figure 2.6B, 2.6C). These results show that the animal integrates the variability of its sensory environment over a time window consisting of the previous 5-30 minutes and uses that information to drive local search behavior .

2.2.3 ASI and ASK sensory neurons are specialized for evaluating variance

Three pairs of sensory neurons (ASI, ASK and AWC) project their dendrites to the nose of the animal where they detect changes in food signals and drive local search behavior (Gray et al., 2005; Wakabayashi et al., 2004) (Figure 2.7A). We examined the role of these neurons in evaluating the size (and thus variance) of the bacterial patch. Interestingly, blocking neurotransmitter release by expressing tetanus toxin light chain fragment (TeTx) (Schiavo et al., 1992) in ASI and ASK but not AWC disrupted learning. AWC-ablated animals also distinguish between the two patch




Figure 2.6. A filter predicts the number of turns during local search. a, A filter representing the contribution of the dispersion in movement orthogonal to the edge predicts that this feature is only important between 5-25 minutes prior to removal from the patch. b, (i, iii) Example tracks on small patches illustrate worms that produce more (top) and fewer (bottom) turns and (ii, iv) comparing predictions from the filter to observed number of turns c, Model prediction provides a good fit of actual search behavior ($r=0.8067$, $p < 10^{-4}$). Data from the example tracks are indicated with a red +. Error bars and shaded regions around the solid lines represent s.e.m. and * indicates significance (t-test, $p < .05$).

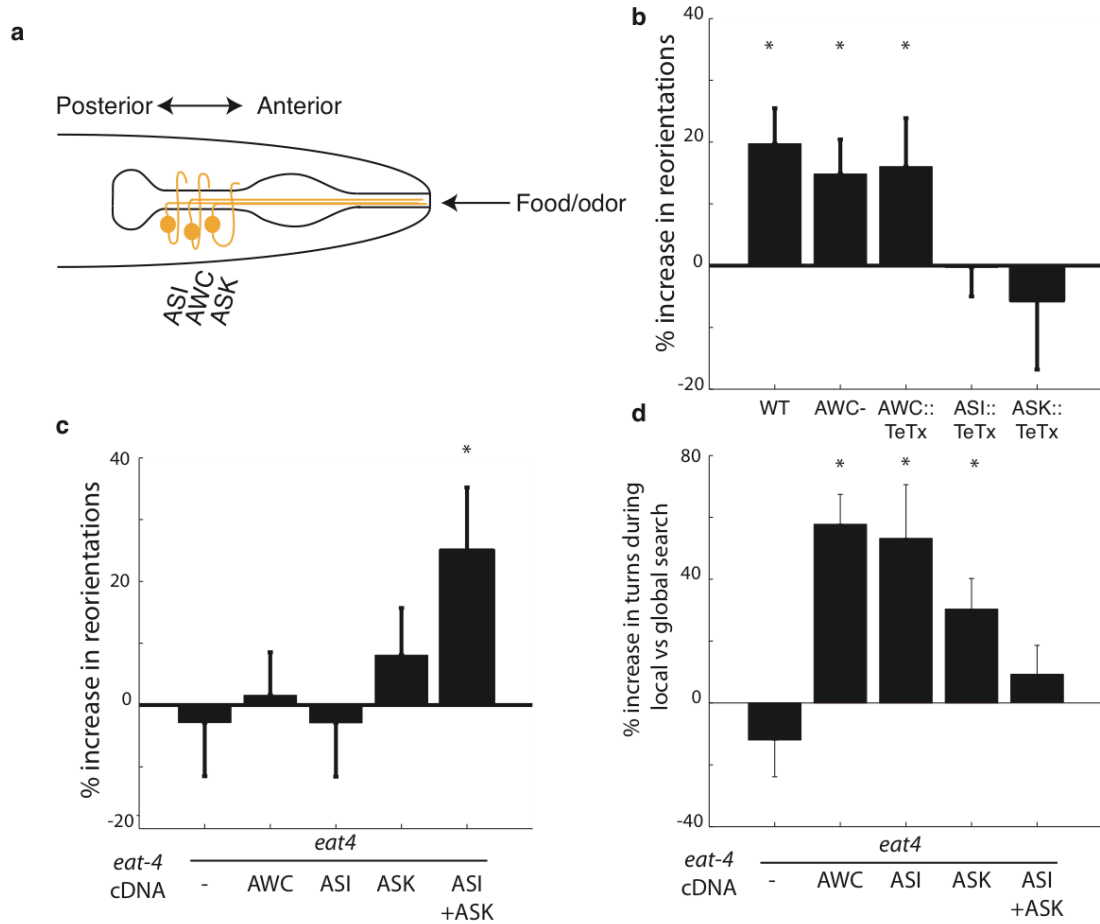


Figure 2.7. Learning requires a unique network. a, Schematic showing the three pairs of amphid sensory neurons that drive local search behavior. b, Blocking neurotransmission by expressing tetanus toxin light chain fragment (TeTx) in ASI and ASK, but not AWC is sufficient to block learning (t-test, $p < 0.05$). c, *eat-4* animals are unable to learn but rescue in both ASI and ASK are sufficient to rescue learning. d, Rescue of *eat-4* in any of the three sensory neurons is sufficient to rescue local search but not learning between patch sizes.

sizes, confirming that the AWC sensory neurons are not required for integrating patch size into local search behavior (Figure 2.7B).

To test whether glutamate release from these neurons was required to drive the food patch modified local search we analyzed the behavior of gene mutants. EAT-4 is the *C. elegans* homolog of the vesicular glutamate transporter and mutants in this gene have severe defects in multiple behaviors including local search (Chalasani et al., 2007; Lee et al., 1999). We found that restoring EAT-4 to both ASI and ASK neurons rescues the ability of the animal to distinguish between different sized patches of food. No single rescue has any effect on this behavior (Figure 2.7C). However, we find that restoring function of the EAT-4 transporter to AWC or ASI or ASK individually rescues the existence of the local search behavior itself (Figure 2.7D). These results show that glutamate release from any of the three neurons (ASI, ASK and AWC) can drive local search behavior, but release from ASI and ASK is necessary and sufficient to distinguish between lawns of different sizes.

To test how ASI and ASK neurons might directly detect variance in the food stimulus we used calcium imaging. We used a microfluidics device that allows us to trap animals and record neural activity while delivering precisely timed stimuli to the nose (Chalasani et al., 2007; Chronis et al., 2007). We compared neural activity patterns using GCaMP calcium indicators (Tian et al., 2009) expressed in ASI, ASK and AWC to small and large changes in bacterial stimuli, analogous to what would be experienced on the centre and edge of the patch, respectively. Consistent with previous results (Chalasani et al., 2007), we observe that AWC responds to removal of both large and small changes in bacteria (Figure 2.8A, 2.8B). Interestingly, ASI

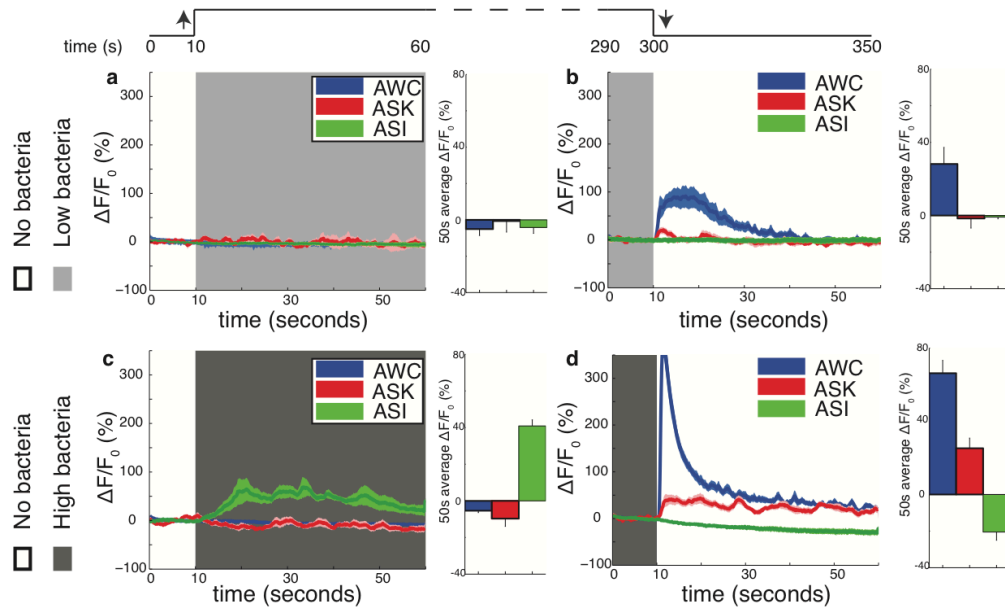


Figure 2.8. Specialized roles of sensory neurons. a-d, Calcium responses in sensory neurons stimulated by small (a,b) and large (c,d) changes in bacterial concentrations. Time course of the imaging experiment is shown with food stimulus added at 10 sec and removed at 300 seconds (arrowheads). Average sensory neuron responses to addition of small (a) or large (c) change and to the removal of small (b) or large (d) change in bacterial stimulus. Average change in fluorescence during the entire 50 sec window after the addition or removal of stimulus is shown for all conditions. Error bars and shaded regions around the solid lines represent s.e.m.

neurons uniquely detect the addition of large, but not small changes in bacterial stimuli (Figure 2.8C, 2.8D). Conversely, ASK sensory neurons respond to removal of large but not small changes in bacteria (Figure 2.8C, 2.8D). Given these activity patterns, it is possible that AWC neurons do not provide sufficient information about the variability the animal experiences at the edge of the patch compared to ASI and ASK. Taken together, these results suggest that ASI and ASK neurons are used together to detect large changes in bacterial concentrations at the edge of a small bacterial patch and release glutamate to signal to downstream neurons and modify local search behavior.

2.2.4 Dopamine suppresses local search behavior

The ASI and ASK sensory neurons synapse onto a common set of interneuron targets including AIY, AIA, AIB and AIZ (White et al., 1986) (Figure 2.9A). To test whether these neurons regulate learning, we blocked neurotransmitter release in individual cells as described above. We found that all of these interneurons played a crucial role such that blocking neurotransmitter release from any of these neurons prevented learning (Figure 2.9B). This suggests that learning is a property of the entire circuit, and not the responsibility of a single neuron or synapse.

We then tested if dopamine, a neuromodulator previously shown to influence learning (Dayan and Balleine, 2002; Hills et al., 2004; Waddell, 2010; Wise, 2004), also modifies search behavior. We found that animals with excessive dopamine (*dat-1*, dopamine transporter) or no dopamine (*cat-2*, tyrosine hydroxylase) at their synapses

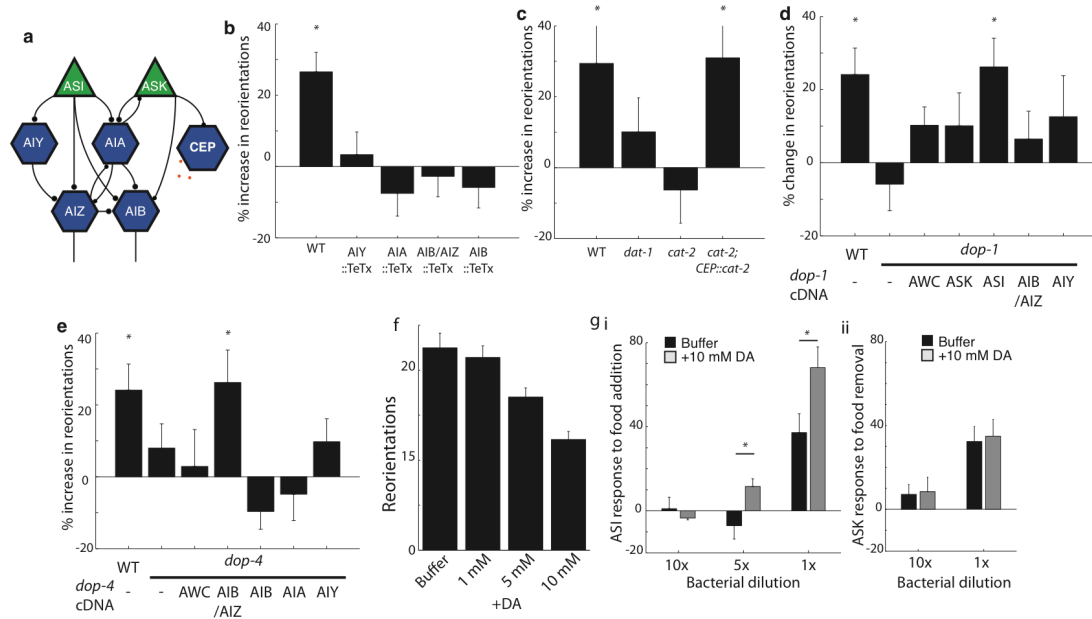


Figure 2.9. Interneuron network for learning. a, Schematic showing the post-synaptic targets of the ASI and ASK sensory neurons required for learning. b, Expressing TeTx in any of the four interneurons blocks learning. c, Learning is also abolished in *cat-2* and *dat-1* mutants, and *cat-2* is specifically required in CEP neurons. d, Mutants in the D1-like dopamine receptor, *dop-1*, are defective in learning and this receptor is specifically required in the ASI sensory neurons. e, *dop-4* acts in AIB and AIZ, interneurons that are downstream of the *dop-1* sensory neurons. f, Addition of dopamine can recapitulate change in search strategy. g, 30s average calcium imaging responses to bacteria is modulated by dopamine in ASI neurons (i) but not ASK neurons (ii).

(Chase and Koelle, 2007) cannot learn. We also found that irrespective of their prior food environment *cat-2* animals behaved similarly to animals removed from large patches. Analogously, *dat-1* mutants behaved similarly to animals removed from small patches. These data suggest that large patch behavior is the default strategy whereas small patch behavior requires dopamine. Interestingly, we are able to restore learning to *cat-2* mutants by expressing CAT-2 specifically in the CEP dopaminergic neurons postsynaptic to the ASK sensory neurons (White et al., 1986) (Figure 2.9C). However, CEP neurons can also directly detect bacteria and drive locomotory behaviors (Sawin et al., 2000). Therefore, we tested whether the sensory function of CEP neurons is required for this learning behavior. We found that *cat-6* mutants, which lack CEP sensory cilia (Perkins et al., 1986), can still learn to distinguish between small and large patches (Figure 2.10). Taken together, these results suggest that CEP neurons might release dopamine acting downstream of the ASK neurons, which respond to removal of large concentrations of food.

We next investigated which cells dopamine is targeting. The *C. elegans* genome encodes several homologs of the mammalian D1- or D2-like receptors (Chase and Koelle, 2007). We found that mutants in the two D1-like receptors, *dop-1* and *dop-4*, are defective in learning (Figure 2.9D, 2.9E). We also observed an antagonistic relationship between the D1-like DOP-1 and the D2-like DOP-3 receptors such that loss of either one prevented learning but loss of both had no effect, consistent with previous results in behavioral paradigms that did not require learning (Chase et al., 2004) (Figure 2.10). However, mutants in other dopamine receptors were able to distinguish between patches of different sizes suggesting that they were not required

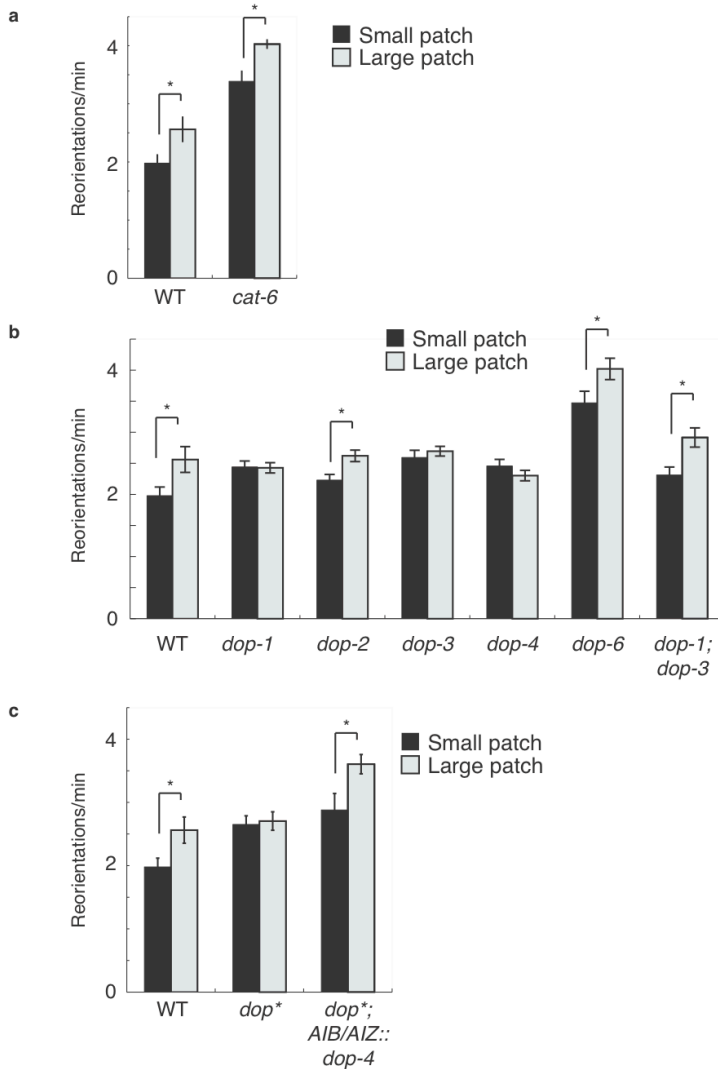


Figure 2.10. Dopamine. a, Among individual dopamine receptor mutants, only *dop-2* and *dop-6* mutants can still learn. *dop-1; dop-3* double mutants regain the ability to learn. *dop-1; dop-2; dop-3; dop-4* quadruple mutants (*dop**) are unable to learn but this phenotype is rescued when *dop-4* is specifically restored to AIB and AIZ. b, Exogenous dopamine decreases turning rate in both wild-type and *crh-1* mutants. Error bars represent s.e.m. and * indicates significance (t-test, $p < .05$).

for this learning (Figure 2.10). Using a cell-specific rescue approach, we have found that DOP-1 receptors are required in ASI sensory neurons (Figure 2.9D). Similarly, the DOP-4 receptors are required in both the AIB and AIZ interneurons, which are postsynaptic to ASI (White et al., 1986) (Figure 2.9E). To test whether the effect of dopamine was specific to this neural circuit, we examined the *dop-1;dop-2;dop-3;dop-4 (dop*)* quadruple mutant (Gaglia and Kenyon, 2009). This mutant is unable to discriminate patch size, though learning is restored when DOP-4 receptor function is rescued in AIB and AIZ interneurons alone (Figure 2.10). Taken together, these results show that dopamine released by CEP neurons is sensed by D1-like receptors on ASI sensory and AIB and AIZ interneurons to regulate learning.

To test whether dopamine was required during learning or during local search we manipulated the level of this neurotransmitter exogenously. We found that adding dopamine to the plate containing food was sufficient to modify local search behavior. Given sufficient dopamine, animals on these plates behaved similarly to animals removed from small patches of bacteria (Figure 2.9). However, adding dopamine to the plate where the animal performed local search had no effect on the behavior (Figure 2.10). These data suggest that dopamine is required during the learning phase and exogenous dopamine decreases the number of turns the animal will emit.

Next, we tested whether this dopamine can directly modify the target ASI neurons. We found that ASI responses to the addition of large concentrations of food are greatly increased in the presence of dopamine. Consistent with our genetic analysis, we found that ASK neurons are not modified by exogenous dopamine (Figure 2.9). These results indicate that animals exploring a small patch are likely to accumulate

dopamine, which amplifies ASI activity. Increased glutamate release from ASI neurons suppresses turns during local behavior.

2.2.5 CREB influences acquisition time

We then tested whether this learning requires protein synthesis and CREB signaling. Cycloheximide has previously been used in *C. elegans* to block protein synthesis in a dose-dependent manner (Kauffman et al., 2010; Szewczyk et al., 2002). Interestingly, blocking protein synthesis with cycloheximide during the 1-hour small patch exploration prevents this learned behavior (Figure 2.11). Cycloheximide does not have a general effect on behavior as treating animals moved between similar sized patches does not cause any significant change in reorientations (Figure 2.12). Moreover, drug treated animals behave similar to animals removed from large lawns, again confirming that the default search behavior is that seen after removal from large lawns while small lawn is learned.

CREB signaling has previously been shown to be required for long-term memory in a number of organisms including *C. elegans* (Kauffman et al., 2010; Silva et al., 1998). Not surprisingly, we find that *crh-1* mutants are unable to learn in our paradigm. This defect is rescued when CREB function is restored to AIB and AIZ interneurons (Figure 2.11). These data show that CREB functions in AIB and AIZ interneurons to regulate learning of patch sizes.

To understand how CREB influences learning, we over expressed *crh-1* in AIB and AIZ interneurons in wildtype animals. These were allowed to grow overnight on

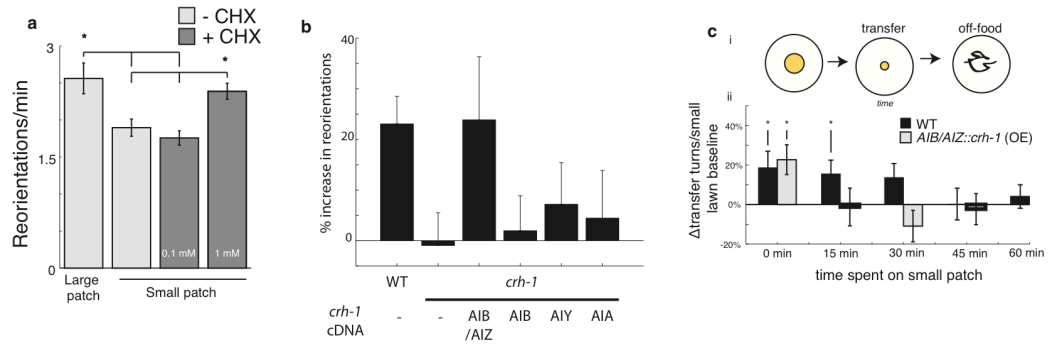


Figure 2.11. Molecular mechanisms of learning. **a**, Inhibition of protein synthesis via cycloheximide (CHX) inhibits learning. **b**, Learning requires *crh-1* in the AIB and AIZ interneurons. **c**, (i) Schematic of transfer assay indicates that (ii) whereas wildtype animals take 30 minutes to adapt to a new environment, animals overexpressing *crh-1* learn much more quickly.

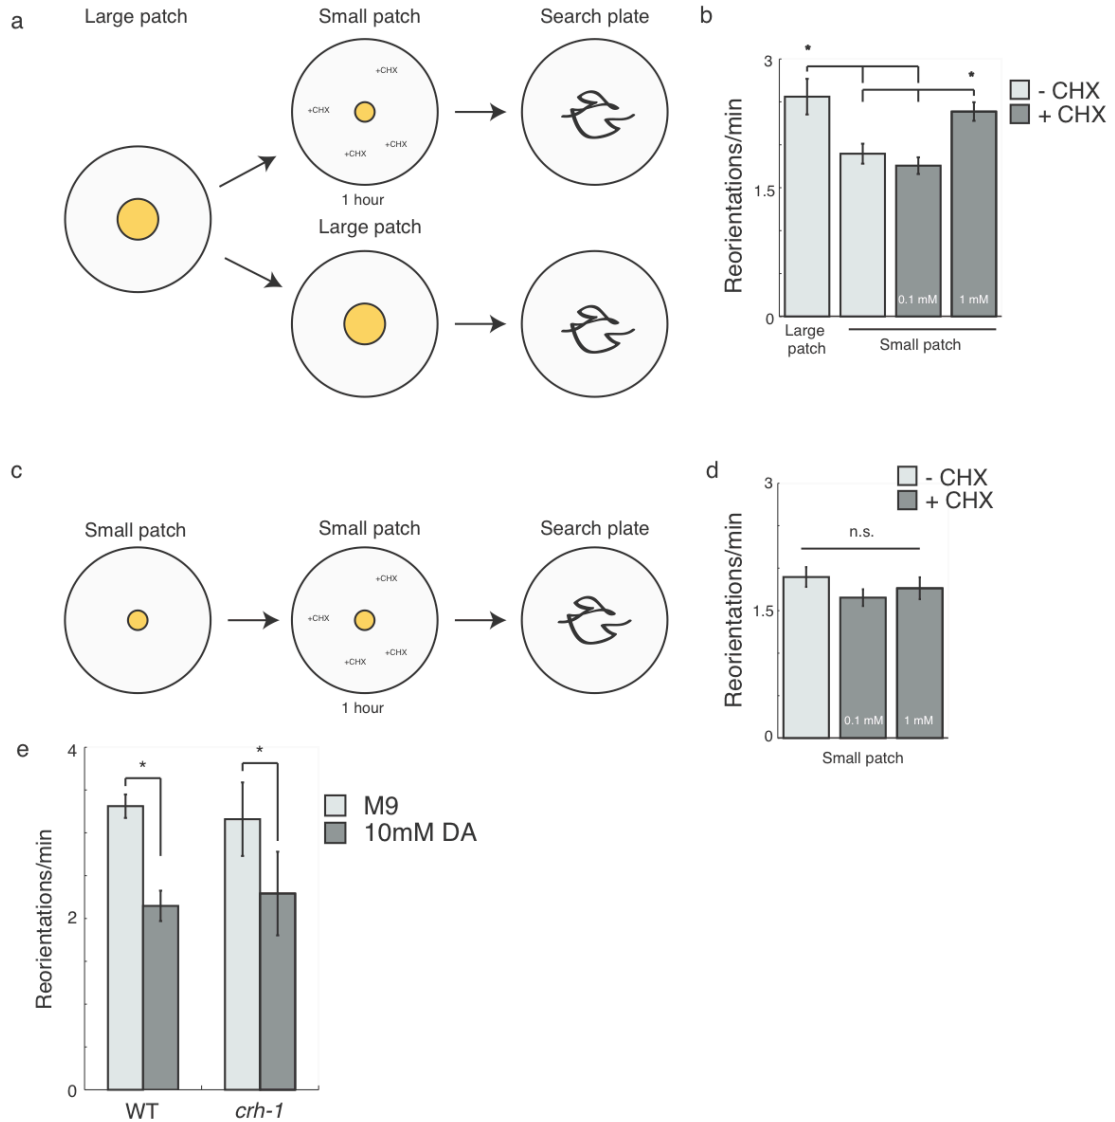


Figure 2.12. Protein synthesis. a, Animals are transferred from large to small lawns containing cycloheximide for 1 hour. b, Cycloheximide prevents learning of small lawn search behavior. c-d, Cycloheximide has no effect on animals transferred between lawns of similar sizes. Error bars represent s.e.m. and * indicates significance (t-test, $p < 0.05$, corrected for multiple comparisons with FDR).

large patches and then transferred to small patches for differing periods of time before their off-food search behavior was assessed. Consistent with the prediction of the sensory filter (Figure 2.10C), we find that wildtype animals take roughly 30 minutes to learn the new environment and switch behavior (Figure 2.10C). In contrast, animals over-expressing CREB need fewer than 15 minutes to learn and execute the new behavior (Figure 2.10C). This suggests that the amount of CREB protein in AIB and AIZ interneurons regulates the time required to acquire new information, a potential temporal mechanism for how CREB influences learning.

2.3 Discussion

The neural circuit driving local search behavior includes neurons that evaluate the variability in the spatial distribution of food in order to modify a complex and long lasting behavioral sequence when animals are removed from food (Gray et al., 2005; Wakabayashi et al., 2004). Learning the unreliability of the environment requires two signals: one representing decreasing bacteria, which is signaled through ASK, and another representing increasing bacteria, signaled through ASI. Our results suggest that this circuit is a subset of the larger behavioral circuit responsible for search (Gray et al., 2005; Wakabayashi et al., 2004) and typically exists in a default configuration arising from low activity in ASI and ASK neurons (Figure 2.13A). However, when ASK, and ASI sensory neurons receive inputs that indicate a varying environment with a less predictable reward, their activity is greatly enhanced, promoting dopamine release from CEP neurons. Dopamine then acts on D1-like receptors on ASI sensory

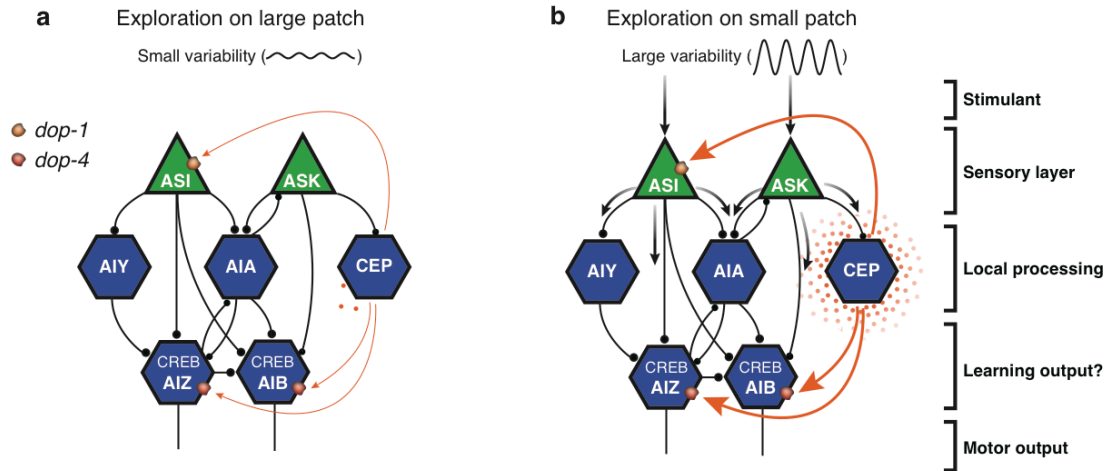


Figure 2.13. Schematic of network for learning. a, The network is relatively quiescent during observations of low variability environments. However, b, during large variability environments dopamine acts on D1-like dopamine receptors and learning occurs through CREB.

and the downstream AIB and AIZ interneurons, in order to modify behavior when confronted with the decision of how to search in a new food-free environment (Figure 2.13B).

2.3.1 Environmental variance modifies local search

Our results show that *C. elegans* evaluates the distribution of food and uses that information to drive a minute-long search strategy when removed from that food. These results are consistent with previous reports showing that animals alter their locomotory patterns and egg-laying when exposed to different foods (Shtonda and Avery, 2006; Waggoner et al., 1998). In particular, animals are more likely to spend more time exploring a patch of ‘bad’ food when compared to ‘good’ food (Shtonda and Avery, 2006). This increased exploration likely reduces the probability of an animal exploring the edge thus reducing environmental variance. Consistently we observe that animals removed from bad foods, like *Bacillus megaterium* (Shtonda and Avery, 2006) experience less variance and turn more frequently when removed from food.

Previous studies have implied that discrete circuit modules process cognitive functions (Friston and Price, 2011). However, the *C. elegans* search circuit contains within itself an ability to modify its own function rather than requiring separate modules for learning. Further, the number of layers, each with multiple neurons, suggests that the computation that is being performed is complex. We show that this circuit is initiated by ASI and ASK neurons that detect variability in the local sensory

environment. Surprisingly, these neurons respond to large, but not small changes in sensory stimuli possibly acting as high-pass filters to maximize information transmission for behaviorally-relevant noise (van Hateren, 1992). Interestingly, the AWC sensory neurons also respond to food (Chalasani et al., 2007), but these neurons detect small fluctuations. In contrast, ASI and ASK neurons only signal large fluctuations that are necessary to estimate environmental variance. We also show that all of three of these sensory neurons (ASI, ASK and AWC) release the neurotransmitter glutamate and drive local search behaviors. Consistent with the imaging data, glutamate release from ASI and ASK, but not AWC is required for learning. These results indicate that this search circuit has evolved to detect changes above a certain threshold and store that information in the downstream neural circuitry over longer periods of time. Interestingly, there exist neurons in the primate brain that also report reward uncertainty (Monosov and Hikosaka, 2013 I think) though it is not known how this value is learned.

2.3.2 Dopamine and CREB signaling interpret sensory information

Interestingly, we find that increased dopamine release resulting from a frequently changing sensory environment achieves plasticity in this learning circuit. This plasticity, in turn, allows the individual to increase risky exploration of a new environment by reducing the number of turns during local search. Our results are consistent with those observed in vertebrate models where dopamine plays a crucial role in motivation, reward and risk-associated behaviors (Schultz, 2002; Schultz et al.,

1997). In particular, increased dopamine has been associated with unpredictable environments and an increase in exploration, while reduced dopamine favors exploitation of previous reward and conserving energy (Beeler, 2012; Beeler et al., 2012). We show that when animals are exposed to large changes in their sensory environment, ASI and ASK sensory neurons are activated leading to increased dopamine release from the CEP neurons and increased exploration (fewer turns) during local search.

We observe this neural circuit modulates the number of turns during local search by integrating the dopamine levels accumulated during a 30-minute exploration on food. This dopamine is then sensed by distinct D-1 like receptors at multiple levels of the circuit, on both ASI sensory and AIB and AIZ interneurons. We suggest that the downstream neural circuitry integrates dopamine signaling via D1 receptors (Chase and Koelle, 2007) and glutamate via AMPA receptors (Chalasani et al., 2007) to drive turn frequencies during local search. Integrating dopamine and glutamate signaling has been shown to play a crucial role in the striatum where it serves to produce long-term changes in synaptic efficiency, both LTP and LTD (Calabresi et al., 1997; Kotter, 1994). These results suggest that our search circuit integrates dopamine and glutamate signaling to modify behaviors that lasts many minutes.

Our results also suggest a novel role for CREB signaling, a pathway that has previously been shown to play a crucial role in regulating learning and memory (Frank and Greenberg, 1994; Silva et al., 1998). Different aspects of learning, such as the value to be learned, learning rate and time scale of learning, have been shown to be under the control of distinct neuromodulatory systems (Doya, 2002). We have

similarly found that learning has multiple independent aspects, including acquisition value and acquisition rate, which are controlled by independent pathways. Dopamine signaling may represent acquisition value, which is indicated as number of turns during local search. In contrast, the amount of CREB protein influences the acquisition rate (time). It is possible that serotonin or other neuromodulators could control the amount of CREB protein in these interneurons which would be consistent with studies showing that serotonin regulates the amount of CREB protein (Lee et al., 2007). Taken together, our results indicate that CREB and dopamine act in two parallel pathways representing information in the circuit.

2.3.3 Two circuit motifs estimate uncertainty

Mechanisms for evaluating environmental changes and variance have been identified in a number of systems. These frequently rely on evaluating local variance in a perceptual receptive field or temporal variance over short timescales (Fairhall et al., 2001; Marr and Hildreth, 1980). However, animals have to evaluate variance over longer timescales as observed in studies involving reward reliability (Monosov and Hikosaka, 2013), but no circuit mechanisms have been proposed. We suggest that two circuit motifs within the local search network combine to estimate uncertainty. The first motif consists of identified sensory neurons that only respond to large changes in the sensory environment, similar to ON and OFF cells in the retina (Wassle, 2004). Just as ON and OFF cells act as edge detectors responding to large increases or decreases in light intensity (Schiller et al., 1986), ASI and ASK respond to large

increases and decreases in bacterial concentration. These results suggest that diverse sensory circuits use parallel ON and OFF pathways to evaluate variance, as a general mechanism. The second motif includes interneurons which are reminiscent of networks proposed for reinforcement learning of average value (Schultz, 1997). This is typically proposed as an error signal sent through dopamine (Doya, 2002; Schultz, 1997; Schultz et al., 1997), similar to the response of ASK through CEP.

We propose that combining circuit motifs may be an evolutionarily efficient method for producing desired behaviors. If the structure of neural circuits have evolved to optimally extract relevant information, this circuit and its computations are likely to be present in other nervous systems. Further studies of this circuit should generate a better understanding of how neural circuits estimate and use uncertain information to drive behaviors, a crucial component of risk strategy and decision-making.

2.4 Experimental procedures

2.4.1 Strains and transgenics

Caenorhabditis elegans strains were maintained as described (Brenner, 1974). A complete list of all strains and transgenics is included in Supplementary Table 2.

cDNAs corresponding to the entire coding sequences of the following genes were amplified using the primers shown below

dop-1

forward 5' TTATGCTAGCATGAACGATTTGCAATGGCCATTG and

reverse 5' TTATGCGGCCGCCTATTCCGGAATGGTTTCCTCG

dop-4

forward 5' AATGTTGGCTTACGGGTCTG and

reverse 5' GCACGTTCTAGTGCAGACCA

crh-1

forward 5' TTATGCTAGCATGGCCACAATGGCGAGCACCTC and

reverse 5' TTATGCGGCCGCTCACATTCCGTCCTTTTCCTTTTCG),

cat-2

forward 5' TTATGGTACCATGTCGTCACTAACCAACAATAC and

reverse 5' TTATGGTACCTCACATTGTAATCGATATTTTC).

Cell-selective expression was achieved using the following promoters: *ceh-36** or *odr-3* for AWC, *sra-9* for ASK, *str-3* for ASI, *odr-2b3a* for AIB and AIZ, *inx-1* for AIB, *ins-1* for AIA, *ttx-3* for AIY, and *p27* for CEP (Bendena et al., 2008; Colon-Ramos et al., 2007; Etchberger et al., 2009; Lin et al., 2010; Roayaie et al., 1998; Troemel et al., 1995; Wenick and Hobert, 2004). For rescue experiments, plasmids at concentrations between 50 and 100 ng/ μ l were microinjected along with 10 ng/ μ l *elt-2::gfp* as a co-injection marker to obtain transgenics using standard protocols (Mello et al., 1991).

2.4.2 Learning assay

Specific volumes [10 μ l (small), 50 μ l, 100 μ l (large), or 200 μ l] of a bacterial culture (OD600 = 0.4) were seeded on NGM agar plates to obtain patches of varying

sizes. Five to six L4 stage animals were allowed to explore these patches overnight. For testing, worms were removed from these plates to a bacteria-free plate and allowed to clean themselves of the remaining bacteria by moving no less than two body lengths, and for no more than 90 seconds. These animals were then transferred to observation plates where they were corralled by a filter paper soaked in 200mM Cu(II)SO₄ solution. Animal behavior on observation plates was recorded for 30 minutes using a Pixelink CCD camera and analyzed by custom software. Data presented was collected from at least 6 plates tested on different days.

Transfer assays were performed similarly. After overnight exploration on large (100μl) patches, animals were transferred to another plate containing either a (100μl) large or (10μl) small patch of bacteria for the indicated amount of time. For assays utilizing exogenous chemicals (dopamine, cycloheximide), the indicated equivalent concentration of chemical was spread on agar plates and allowed to absorb into NGM agar plates for 90 minutes. Bacterial culture (100μl or 10μl) was then seeded on these plates for sixty minutes and allowed to dry before animals were placed on the patch for the given length of time. In certain assays, bacteria were killed in a water bath at 65C for 3hours before plating.

2.4.3 Calcium imaging

A PDMS based microfluidic device was used to record activity from specific neurons expressing GCaMP calcium indicators as described (Chalasani et al., 2007). Salt levels in LB media were found to interfere with neural responses so bacteria had

to be resuspended in buffer solution. We grew a bacterial culture to an OD600 of 0.4 and centrifuged it at 3000 RPM for 10 minutes. To generate bacterial stimuli, the resulting pellet was either diluted 1:100 or left undiluted to mimic small and large changes in bacteria, respectively. GCaMP imaging was performed on a Zeiss inverted microscope using a Photometrics EMCCD camera. Images were captured in Metamorph imaging software at 10fps and analyzed offline using custom code. Baseline F0 was measured as average intensity across the first 1-4 seconds of recording. The ratio of change in fluorescence to the baseline F0 is plotted in our results (Chalasani et al., 2007).

2.4.4 Imaging bacteria patches

Bacteria expressing green fluorescent protein (Labrousse et al., 2000) were plated identically to bacteria used for learning assays. Five plates each containing patches of 10 μ l, 50 μ l, 100 μ l, and 200 μ l bacteria were imaged on a Zeiss Stereo microscope using a Zeiss MRM CCD camera. As the highest fluorescence was at the edge of the patch, peaks were fit to an ellipse in order to extract fluorescence profiles and the bacterial centre. For each individual patch, a profile was extracted every 36 degrees and averaged. Normalized profiles were found by dividing the distance from the edge by the radius of the patch, so that the distance at the edge was 0 and at the centre was 1. Profiles are well fit to a 1/R distribution from the edge to the centre of the patch. We chose the 20% of the patch closest to where the bacteria ends on a small patch as the 'edge' of the patch and the other 80% to represent the centre. The edge of

larger patches was chosen to represent the same absolute distance from the transition between bacteria and no bacteria.

2.4.5 Simulations

A random walk simulation was written in MATLAB where forward movement was maintained. An arbitrary number of turns were inserted into the path to change the orientation at random intervals. “Omega” turns have a mean of 135 degrees and “non-omega” turns included motifs with a mean 60 degrees. 100 iterations of the simulation were run for each number of turns, and the mean-squared diffusion calculated from these simulations. Mean-squared diffusion was defined as the square of the distance travelled across an arbitrary time period. To analyze worm behavioral data, tracks were binned into one-minute intervals and the number of pirouettes that began in the bin was summed. The mean-squared diffusion was the square of the Euclidean distance between the initial position and the ending position in the bin.

2.4.6 Prediction and Maximum Noise Entropy

In order to identify the variables responsible for on-food learning, we grew L4 animals overnight on large bacterial patches. The next day, single young adults were transferred to small patches and their movements recorded for one hour. Immediately afterward, animals were moved to an observation plate containing a copper ring and recorded as described in the learning assay above. Movement on-food was analyzed, and body and nose position extracted. Initial comparisons between on-food variables

and off-food turning showed several variables that weakly predicted turning behavior: time on edge of patch, variability in position (total movement), and mean distance from edge. The strongest predictor was the dispersion of time spent on the edge, or amount of movement on and off the edge towards either the centre of the patch or away from the food. Notably, using the position of the nose gave a significantly better prediction than using the centre of mass of the body (data not shown). Other variables, such as time spent off food and variance in position while not on the edge (the variance in the set of all the positions excluding time on edge), showed no relevance to behavior (Figure 2.5).

Filters were extracted via the method of Maximum Noise Entropy (MNE) (Fitzgerald et al., 2011). Whereas methods such as reverse correlation will find the input that best correlates with the output, MNE finds the filter that will maximize the uncertainty across trials in order to be consistent with the known input/output relationships while making no assumptions about anything else. Here, we binned parameters into 3-minute time periods in order to find the time bins that are most important in determining the behavioral output, the number of turns upon removal from food. The MNE feature was identified by maximizing the noise entropy, or equivalently minimizing mutual information between a vector f that will produce $f * (\text{input}) = \text{turns}$. MNE filters were extracted in four jack-knives which were then averaged. Final turning was predicted using the stimulus energy from the MNE filter and a subset of the data was then fitted with nonlinearity in accordance with the linear-nonlinear model (Fitzgerald et al., 2011).

Table 2.1. Strain list.

Strain	Genotype	Name
N2	Bristol strain	WT
IV237	ueEx8 [ins-1::GCaMP3, unc-122::gfp]; oyIs [ceh-36del::caspase-3(p12)::nz, ceh- 36del::cz::caspase-3(p17), srtx-1::gfp, unc- 122::dsRED]	AWC::Caspase
	<i>nm3315</i>	crh-1
IV216	<i>ueEx131 [odr-3::TeTx::mCherry; elt2::GFP]</i>	AWC::TeTx
IV205	<i>ueEx122 [str-3::TeTx::GFP; elt2::sl2GFP]</i>	ASI::TeTx
CX11576	<i>kyEx3097 [sra-9::TNT::mCherry; elt2::GFP]</i>	ASK::TeTx
IV217	<i>ueEx132 [ttx-3::TeTx::mCherry; elt2::sl2GFP]</i>	AIY::TeTx
IV203	<i>ueEx120 [ins-1::TeTx::mCherry; elt2::sl2GFP]</i>	AIA::TeTx
IV316	<i>ueEx196 [inx-1::TeTx::GFP; elt2::sl2GFP]</i>	AIB/AIZ::TeTx
IV314	<i>ueEx194 [odr2b3a::TeTx::GFP; elt2::sl2GFP]</i>	AIB::TeTx
CX10536	<i>kyEx2595 [str-2::gcamp2.2b, unc-122::gfp]</i>	AWC (imaging)
CX10979	<i>kyEx2865 [sra-6::gcamp3; ofm-1::gfp]</i>	ASI (imaging)
CX10981	<i>kyEx2866 [sra-9::gcamp2.2b; ofm-1::gfp]</i>	ASK (imaging)
CB1112	<i>cat-2(e1112) II</i>	<i>cat-2</i>
RM2702	<i>dat-1(ok157) III</i>	<i>dat-1</i>
LX645	<i>dop-1(vs100) V</i>	<i>dop-1</i>
LX702	<i>dop-2(vs105) V</i>	<i>dop-2</i>
LX703	<i>dop-3(vs106) X</i>	<i>dop-3</i>
FG58	<i>dop-4(tm1392) X</i>	<i>dop-4</i>
CF2805	<i>dop-1(vs100); dop-2(vs105); dop-3(vs106); dop-4(ok1321)</i>	<i>dop*</i>
IV111	<i>cat-2(e1112); ueEx51 [p27::cat-2::GFP; elt2::GFP]</i>	cat-2; CEP::cat-2
IV83	<i>dop-1(vs100); ueEx35 [ceh36*::dop-1::GFP; elt2::GFP]</i>	dop-1; AWC::dop-1
IV377	<i>dop-1(vs100); ueEx246 [sra-9::dop-1; elt2::GFP]</i>	dop-1; ASK::dop-1
IV376	<i>dop-1(vs100); ueEx245 [str-3::dop-1; elt2::GFP]</i>	dop-1; ASI::dop-1
IV86	<i>dop-1(vs100); ueEx38 [odr2b3a::dop-1::GFP; elt2::GFP]</i>	dop-1; AIB/AIZ::dop-1

Table 2.1. Strain list, continued.

Strain	Genotype	Name
IV88	<i>dop-1(vs100); ueEx40 [ttx3::<dop-1>::GFP; elt2::<gfp]< i=""></gfp]<></dop-1></i>	dop-1; AIY:: <dop-1< td=""> </dop-1<>
IV90	<i>dop-4(tm1392); ueEx42 [ceh36*::<dop-4>::GFP; elt2::<gfp]< i=""></gfp]<></dop-4></i>	dop-4; AWC:: <dop-4< td=""> </dop-4<>
IV48	<i>dop-4(tm1392); ueEx25 [odr2b3a::<dop-4>::GFP; elt2::<gfp]< i=""></gfp]<></dop-4></i>	dop-4; AIB/AIZ:: <dop-4< td=""> </dop-4<>
IV208	<i>dop-4(tm1392); ueEx125 [inx-1::<dop-4>::GFP; elt2::<gfp]< i=""></gfp]<></dop-4></i>	dop-4; AIB:: <dop-4< td=""> </dop-4<>
IV46	<i>dop-4(tm1392); ueEx23 [ins-1::<dop-4>::GFP; elt2::<gfp]< i=""></gfp]<></dop-4></i>	dop-4; AIA:: <dop-4< td=""> </dop-4<>
IV49	<i>dop-4(tm1392); ueEx26 [ttx-3::<dop4>::GFP; elt2::<gfp]< i=""></gfp]<></dop4></i>	dop-4; AIY:: <dop-4< td=""> </dop-4<>
IV296	<i>dop-1(vs100); dop-2(vs105); dop-3(vs106); dop-4(ok1321); ueEx186 [odr2b3a::<dop-4>::GFP; elt2::<gfp]< i=""></gfp]<></dop-4></i>	dop*; AIB/AIZ:: <dop-4< td=""> </dop-4<>
IV85	<i>crh-1(nn3315); ueEx37 [odr2b3a::<crh-1>::GFP; elt2::<gfp]< i=""></gfp]<></crh-1></i>	crh-1; AIB/AIZ:: <crh-1< td=""> </crh-1<>
IV204	<i>crh-1(nn3315); ueEx129 [inx-1::<crh-1>::GFP; elt2::<gfp]< i=""></gfp]<></crh-1></i>	crh-1; AIB:: <crh-1< td=""> </crh-1<>
IV84	<i>crh-1(nn3315); ueEx36 [ttx-3::<crh-1>::GFP; elt2::<gfp]< i=""></gfp]<></crh-1></i>	crh-1; AIY:: <crh-1< td=""> </crh-1<>
IV319	<i>crh-1(nn3315); ueEx199 [ins-1::<crh-1>::GFP; elt2::<gfp]< i=""></gfp]<></crh-1></i>	crh-1; AIA:: <crh-1< td=""> </crh-1<>
IV368	<i>ueEx238 [odr2b3a::<crh-1>::GFP; elt2::<gfp]< i=""></gfp]<></crh-1></i>	AIB/AIZ:: <crh-1 (oe)<="" td=""> </crh-1>

Table 2.2. Search numbers

Strain	Small patch, local search	Large patch, local search
AWC::TeTx	2.7045 +/- 0.1548	3.2485 +/- 0.1269
AWC(-)	2.5162 +/- 0.1186	2.9660 +/- 0.1108
ASI::TeTx	3.1104 +/- 0.1063	3.1020 +/- 0.1012
ASK::TeTx	2.2202 +/- 0.2039	2.0947 +/- 0.1999
AIY::TeTx	3.6861 +/- 0.1436	3.8088 +/- 0.1816
AIA::TeTx	5.0386 +/- 0.2544	4.6587 +/- 0.2189
AIB/AIZ::TeTx	2.9408 +/- 0.1110	2.8598 +/- 0.1292
AIB::TeTx	4.0473 +/- 0.2280	3.8090 +/- 0.1175
crh-1	3.2709 +/- 0.1501	3.2402 +/- 0.1500
crh-1;		
AIB/AIZ::crh-1	2.9604 +/- 0.2389	3.6670 +/- 0.2136
crh-1; AIB::crh-1	3.5243 +/- 0.1742	3.5932 +/- 0.1674
crh-1; AIY::crh-1	3.8685 +/- 0.1583	3.9062 +/- 0.2150
crh-1; AIA::crh-1	3.8276 +/- 0.2741	3.9984 +/- 0.2185
cat-2	2.4420 +/- 0.1404	2.2879 +/- 0.1760
dat-1	1.3955 +/- 0.0977	1.6171 +/- 0.1113
cat-2; CEP::cat-2	2.1057 +/- 0.1193	2.8833 +/- 0.1220
dop-1	2.4415 +/- 0.0956	2.4275 +/- 0.0830
dop-2	2.2297 +/- 0.0917	2.6204 +/- 0.0925
dop-3	2.5943 +/- 0.1157	2.6959 +/- 0.0782
dop-4	2.4569 +/- 0.1067	2.3023 +/- 0.0841
dop-6	3.4725 +/- 0.1870	4.0196 +/- 0.1713
dop-1;dop-3	2.3120 +/- 0.1286	2.9153 +/- 0.1541
dop-1; AWC::dop-1	4.7353 +/- 0.2012	5.2203 +/- 0.1748
dop-1; AIB/AIZ::dop-1	3.5335 +/- 0.2162	3.7626 +/- 0.1797
dop-1; AIY::dop-1	3.4900 +/- 0.2595	3.9294 +/- 0.3118
dop-1; ASI::dop-1	3.1153 +/- 0.1605	3.9317 +/- 0.1873
dop-1; ASK::dop-1	1.6776 +/- 0.1347	2.0033 +/- 0.1044
dop-4; AWC::dop-4	2.4442 +/- 0.1750	2.5147 +/- 0.2207
dop-4; AIB/AIZ::dop-4	3.4062 +/- 0.1572	4.3003 +/- 0.2785
dop-4; AIB::dop-4	4.3182 +/- 0.1752	4.0589 +/- 0.2087
dop-4; AIA::dop-4	3.4157 +/- 0.2202	3.5904 +/- 0.1657
dop-4; AIY::dop-4	3.7676 +/- 0.2102	4.2686 +/- 0.1429
dop*	2.6509 +/- 0.1366	2.7051 +/- 0.1463
dop*; AIB/AIZ::dop-4	2.8798 +/- 0.2614	3.6048 +/- 0.1521
WT	3.2518 +/- 0.1057	4.0021 +/- 0.1199

2.5 Acknowledgments

Chapter 2, in full, is a manuscript in revision. Adam J Calhoun, Navin Pokala, James AJ Fitzpatrick, Tatyana O Sharpee, Sreekanth H Chalasani. Circuit mechanisms for evaluating uncertainty in *Caenorhabditis elegans*. In revision, 2014. The dissertation author was the primary investigator and author of this paper.

2.6 References

- Avery, L., and Shtonda, B.B. (2003). Food transport in the *C. elegans* pharynx. *J Exp Biol* 206, 2441-2457.
- Beeler, J.A. (2012). Thorndike's Law 2.0: Dopamine and the Regulation of Thrift. *Front Neurosci* 6, 116.
- Beeler, J.A., Frazier, C.R., and Zhuang, X. (2012). Dopaminergic enhancement of local food-seeking is under global homeostatic control. *Eur J Neurosci* 35, 146-159.
- Bendena, W.G., Boudreau, J.R., Papanicolaou, T., Maltby, M., Tobe, S.S., and Chin-Sang, I.D. (2008). A *Caenorhabditis elegans* allatostatin/galanin-like receptor NPR-9 inhibits local search behavior in response to feeding cues. *Proc Natl Acad Sci U S A* 105, 1339-1342.
- Brenner, S. (1974). The genetics of *Caenorhabditis elegans*. *Genetics* 77, 71-94.
- Calabresi, P., Pisani, A., Centonze, D., and Bernardi, G. (1997). Synaptic plasticity and physiological interactions between dopamine and glutamate in the striatum. *Neurosci Biobehav Rev* 21, 519-523.
- Chalasani, S.H., Chronis, N., Tsunozaki, M., Gray, J.M., Ramot, D., Goodman, M.B., and Bargmann, C.I. (2007). Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*. *Nature* 450, 63-70.
- Chalfie, M., Sulston, J.E., White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1985). The neural circuit for touch sensitivity in *Caenorhabditis elegans*. *J Neurosci* 5, 956-964.
- Chase, D.L., and Koelle, M.R. (2007). Biogenic amine neurotransmitters in *C. elegans*. *WormBook*, 1-15.
- Chase, D.L., Pepper, J.S., and Koelle, M.R. (2004). Mechanism of extrasynaptic dopamine signaling in *Caenorhabditis elegans*. *Nat Neurosci* 7, 1096-1103.
- Chronis, N., Zimmer, M., and Bargmann, C.I. (2007). Microfluidics for in vivo imaging of neuronal and behavioral activity in *Caenorhabditis elegans*. *Nat Methods* 4, 727-731.
- Colon-Ramos, D.A., Margeta, M.A., and Shen, K. (2007). Glia promote local synaptogenesis through UNC-6 (netrin) signaling in *C. elegans*. *Science* 318, 103-106.

- Dayan, P., and Balleine, B.W. (2002). Reward, motivation, and reinforcement learning. *Neuron* 36, 285-298.
- de Bono, M., and Maricq, A.V. (2005). Neuronal substrates of complex behaviors in *C. elegans*. *Annu Rev Neurosci* 28, 451-501.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Netw* 15, 495-506.
- Escobar, C., Martinez-Merlos, M.T., Angeles-Castellanos, M., del Carmen Minana, M., and Buijs, R.M. (2007). Unpredictable feeding schedules unmask a system for daily resetting of behavioural and metabolic food entrainment. *Eur J Neurosci* 26, 2804-2814.
- Etchberger, J.F., Flowers, E.B., Poole, R.J., Bashllari, E., and Hobert, O. (2009). Cis-regulatory mechanisms of left/right asymmetric neuron-subtype specification in *C. elegans*. *Development* 136, 147-160.
- Fairhall, A.L., Lewen, G.D., Bialek, W., and de Ruyter Van Steveninck, R.R. (2001). Efficiency and ambiguity in an adaptive neural code. *Nature* 412, 787-792.
- Fitzgerald, J.D., Rowekamp, R.J., Sincich, L.C., and Sharpee, T.O. (2011). Second order dimensionality reduction using minimum and maximum mutual information models. *PLoS Comput Biol* 7, e1002249.
- Frank, D.A., and Greenberg, M.E. (1994). CREB: a mediator of long-term memory from mollusks to mammals. *Cell* 79, 5-8.
- Friston, K.J., and Price, C.J. (2011). Modules and brain mapping. *Cogn Neuropsychol* 28, 241-250.
- Gaglia, M.M., and Kenyon, C. (2009). Stimulation of movement in a quiescent, hibernation-like form of *Caenorhabditis elegans* by dopamine signaling. *J Neurosci* 29, 7302-7314.
- Gray, J.M., Hill, J.J., and Bargmann, C.I. (2005). A circuit for navigation in *Caenorhabditis elegans*. *Proc Natl Acad Sci U S A* 102, 3184-3191.
- Hills, T., Brockie, P.J., and Maricq, A.V. (2004). Dopamine and glutamate control area-restricted search behavior in *Caenorhabditis elegans*. *J Neurosci* 24, 1217-1225.
- Hills, T.T., Kalff, C., and Wiener, J.M. (2013). Adaptive Levy processes and area-restricted search in human foraging. *PLoS ONE* 8, e60488.

- Kauffman, A.L., Ashraf, J.M., Corces-Zimmerman, M.R., Landis, J.N., and Murphy, C.T. (2010). Insulin signaling and dietary restriction differentially influence the decline of learning and memory with age. *PLoS Biol* 8, e1000372.
- Kotter, R. (1994). Postsynaptic integration of glutamatergic and dopaminergic signals in the striatum. *Prog Neurobiol* 44, 163-196.
- Labrousse, A., Chauvet, S., Couillault, C., Kurz, C.L., and Ewbank, J.J. (2000). *Caenorhabditis elegans* is a model host for *Salmonella typhimurium*. *Curr Biol* 10, 1543-1545.
- Lee, R.Y., Sawin, E.R., Chalfie, M., Horvitz, H.R., and Avery, L. (1999). EAT-4, a homolog of a mammalian sodium-dependent inorganic phosphate cotransporter, is necessary for glutamatergic neurotransmission in *caenorhabditis elegans*. *J Neurosci* 19, 159-167.
- Lee, S.H., Lim, C.S., Park, H., Lee, J.A., Han, J.H., Kim, H., Cheang, Y.H., Lee, Y.S., Ko, H.G., Jang, D.H., et al. (2007). Nuclear translocation of CAM-associated protein activates transcription for long-term facilitation in *Aplysia*. *Cell* 129, 801-812.
- Lin, C.H., Tomioka, M., Pereira, S., Sellings, L., Iino, Y., and van der Kooy, D. (2010). Insulin signaling plays a dual role in *Caenorhabditis elegans* memory acquisition and memory retrieval. *J Neurosci* 30, 8001-8011.
- MacLean, E.L., Mandalaywala, T.M., and Brannon, E.M. (2012). Variance-sensitive choice in lemurs: constancy trumps quantity. *Anim Cogn* 15, 15-25.
- Marr, D., and Hildreth, E. (1980). Theory of edge detection. *Proc R Soc Lond B Biol Sci* 207, 187-217.
- Mello, C.C., Kramer, J.M., Stinchcomb, D., and Ambros, V. (1991). Efficient gene transfer in *C.elegans*: extrachromosomal maintenance and integration of transforming sequences. *Embo J* 10, 3959-3970.
- Monosov, I.E., and Hikosaka, O. (2013). Selective and graded coding of reward uncertainty by neurons in the primate anterodorsal septal region. *Nat Neurosci* 16, 756-762.
- Perkins, L.A., Hedgecock, E.M., Thomson, J.N., and Culotti, J.G. (1986). Mutant sensory cilia in the nematode *Caenorhabditis elegans*. *Dev Biol* 117, 456-487.
- Platt, M.L., and Huettel, S.A. (2008). Risky business: the neuroeconomics of decision making under uncertainty. *Nat Neurosci* 11, 398-403.

- Roayaie, K., Crump, J.G., Sagasti, A., and Bargmann, C.I. (1998). The G alpha protein ODR-3 mediates olfactory and nociceptive function and controls cilium morphogenesis in *C. elegans* olfactory neurons. *Neuron* 20, 55-67.
- Sawin, E.R., Ranganathan, R., and Horvitz, H.R. (2000). *C. elegans* locomotory rate is modulated by the environment through a dopaminergic pathway and by experience through a serotonergic pathway. *Neuron* 26, 619-631.
- Schiavo, G., Benfenati, F., Poulain, B., Rossetto, O., Polverino de Laureto, P., DasGupta, B.R., and Montecucco, C. (1992). Tetanus and botulinum-B neurotoxins block neurotransmitter release by proteolytic cleavage of synaptobrevin. *Nature* 359, 832-835.
- Schiller, P.H., Sandell, J.H., and Maunsell, J.H. (1986). Functions of the ON and OFF channels of the visual system. *Nature* 322, 824-825.
- Schultz, W. (1997). Dopamine neurons and their role in reward mechanisms. *Curr Opin Neurobiol* 7, 191-197.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron* 36, 241-263.
- Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593-1599.
- Shtonda, B.B., and Avery, L. (2006). Dietary choice behavior in *Caenorhabditis elegans*. *J Exp Biol* 209, 89-102.
- Silva, A.J., Kogan, J.H., Frankland, P.W., and Kida, S. (1998). CREB and memory. *Annu Rev Neurosci* 21, 127-148.
- Sommerfeld, J., Kato, A., Ropert-Coudert, Y., Garthe, S., and Hindell, M.A. (2013). Foraging parameters influencing the detection and interpretation of area-restricted search behaviour in marine predators: a case study with the masked booby. *PLoS ONE* 8, e63742.
- Szewczyk, N.J., Peterson, B.K., and Jacobson, L.A. (2002). Activation of Ras and the mitogen-activated protein kinase pathway promotes protein degradation in muscle cells of *Caenorhabditis elegans*. *Mol Cell Biol* 22, 4181-4188.
- Thums, M., Bradshaw, C.J., and Hindelli, M.A. (2011). In situ measures of foraging success and prey encounter reveal marine habitat-dependent search strategies. *Ecology* 92, 1258-1270.
- Tian, L., Hires, S.A., Mao, T., Huber, D., Chiappe, M.E., Chalasani, S.H., Petreanu, L., Akerboom, J., McKinney, S.A., Schreiter, E.R., et al. (2009). Imaging

neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat Methods* 6, 875-881.

Troemel, E.R., Chou, J.H., Dwyer, N.D., Colbert, H.A., and Bargmann, C.I. (1995). Divergent seven transmembrane receptors are candidate chemosensory receptors in *C. elegans*. *Cell* 83, 207-218.

van Hateren, J.H. (1992). A theory of maximizing sensory information. *Biol Cybern* 68, 23-29.

Waddell, S. (2010). Dopamine reveals neural circuit mechanisms of fly memory. *Trends Neurosci* 33, 457-464.

Waggoner, L.E., Zhou, G.T., Schafer, R.W., and Schafer, W.R. (1998). Control of alternative behavioral states by serotonin in *Caenorhabditis elegans*. *Neuron* 21, 203-214.

Wakabayashi, T., Kitagawa, I., and Shingai, R. (2004). Neurons regulating the duration of forward locomotion in *Caenorhabditis elegans*. *Neurosci Res* 50, 103-111.

Wark, B., Fairhall, A., and Rieke, F. (2009). Timescales of inference in visual adaptation. *Neuron* 61, 750-761.

Wassle, H. (2004). Parallel processing in the mammalian retina. *Nat Rev Neurosci* 5, 747-757.

Watson, K.K. (2008). Evolution, risk, and neural representation. *Ann N Y Acad Sci* 1128, 8-12.

Wenick, A.S., and Hobert, O. (2004). Genomic cis-regulatory architecture and trans-acting regulators of a single interneuron-specific gene battery in *C. elegans*. *Dev Cell* 6, 757-770.

White, J.G., Southgate, E., Thomson, J.N., and Brenner, S. (1986). The structure of the nervous system of the nematode *Caenorhabditis elegans*. *Phil. Transact. R. Soc. Lond. B* 314, 1-340.

Wise, R.A. (2004). Dopamine, learning and motivation. *Nat Rev Neurosci* 5, 483-494.

3. Emergent discrete decisions within continuous maximally informative search strategies

3.1 Abstract

How animals explore their environment over areas many orders of magnitude larger than their body size is an unsolved decision-making problem that has broad practical and environmental applications. Experiments show that the off-food search patterns of the nematode *C. elegans* exhibit a stereotypic change in behavior from a high-turn search state lasting many minutes (“local search”) to a “global search” characterized by reduced turning and covering large expanses of their environment. We show that this transition can be quantitatively explained by a maximally informative search strategy where the searcher attempts to continually increase information about the possible food location. While this strategy converges to chemotaxis close to a food source, it offers a number of predictions that are distinct from chemotaxis, some of which we have experimentally verified. Further, we show that the maximally informative search implicitly contains a decision variable that can be approximated by a drift-diffusion model. The diffusion variable reflects the probability that the food is not contained in the local area, and the global search begins when this probability reaches one. The mapping between drift-diffusion and maximally informative models points to simple heuristic computations that can be implemented in the brain to yield near-optimal behavioral performance.

3.2 Introduction

Making appropriate decisions is central to organism's survival. Foraging behaviors of small animals that search for food over area many times their size provide a rich quantitative paradigm within which one can study principles and neural mechanisms of decision-making. Recent experimental studies (Gray et al., 2005; Hills et al., 2004; Wakabayashi et al., 2004; Chalasani et al., 2007) report a striking and robust aspect in the foraging patterns of *C. elegans* once the animals have been moved from plates containing bacterial food to plates without bacteria. During the initial period after removal from food, termed 'local search', the animal performs an intense search around the area where it believes food is likely to be located (Fig. 3.1a). This period is characterized by an increased number of abrupt turns (Fig. 3.1b). After approximately 15 minutes, wild-type animals reduce the number of turns to a basal rate (Fig. 3.1b). This results in more extended trajectories (Fig. 3.1a) and allows the animal to search a much larger area. For these reasons, the second part of the search is termed 'global search.' Although *C. elegans* is traditionally considered to be a chemotactic searcher, moving up or down chemical gradients to find the source of an odorant, in these conditions animals have no chemical gradient to follow. Thus, we set out to explore whether the observed animals' behavior could reflect an active inference process.

An attractive possibility is that animals' behavior may in general be guided by the need to maximize information about how to achieve the behavioral goal. This framework was recently shown to account for distinct properties of human eye movement search (Najemnik and Geisler, 2005) and animals navigation in turbulent

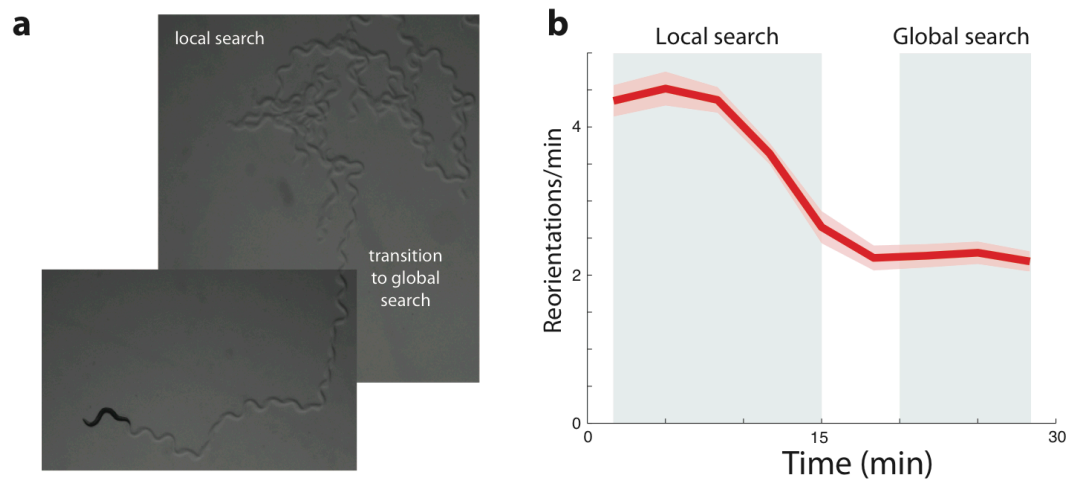


Figure 3.1. *C. elegans* displays transition in search behavior. a, Animals search the local area by emitting a large number of turns before abruptly transitioning to a global search. b, Across many animals, this transition is readily apparent in the mean turning rate.

environment (Vergassola et al., 2007; Masson et al., 2009). Importantly, information maximization strategies analyzed in those contexts have shown a number of features that would not be expected if animals are following a more familiar gradient - following strategy that aims to bring the animal directly along a path to a goal. For example, humans sometimes make saccades to examine a region between, instead of directly at, two likely locations for a target (Najemnik et al., 2005). Birds and moths zigzag across wind currents to find the source of a plume (Vergassola et al., 2007). Both of these features can be accounted by the adaptation of a maximally informative search strategy to the appropriate behavioral context, and would not be expected during a gradient-like search. At the same time, for positions closer to the target, maximal information solutions converge to chemotaxis (Vergassola et al., 2007), and therefore can be viewed as a generalization of this classical approach to animal behavior. An important difference however between the above examples of animal behavior and the case of *C. elegans* is that following removal from food worms are not receiving any positive sensory cues. Even in the case of turbulent environment the search cannot proceed successfully if odorants are not detected at least from time to time. Thus, it is important to determine what kind of search patterns one would expect to observe for animals deprived of sensory inputs if they are following an optimal maximally informative strategy.

3.3 Results

This work attempts to determine whether an information-maximization strategy is consistent with animal behavior. To do this, we begin by modeling the prior belief about the distribution of food as a two-dimensional Gaussian. We then compute maximally informative trajectories where the action at each time step is selected to maximize information about the location of the food source. During search, possible steps are evaluated based on estimated probabilities to detect an odorant given the probability distribution for the location of the source $P(\vec{r})$, see Materials and methods for details. Importantly, the searcher gains information both from odorant detection and non-detection events. We find that these maximally informative solutions exhibit an abrupt transition between what is at least conceptually consistent with a local and global search (Fig. 3.2a). The initial stage of the search consists of large number of turns is localized within the extent of the initial prior distribution. As the search progress the distribution of food source locations $P(\vec{r})$ is continuously eroded, because no odorants are detected in the regions with initially high likelihood for food source. By the time the animals trajectory straightens going directly to the region boundary, no substantial peaks remain in the distribution (Fig. 3.2b).

A critical aspect of the model is that the overall probability that the source of food is located within the modeled area A (the full extent of the area shown in Fig. 3.2),

$$p_t(A) = \int [d\vec{r} P_t(\vec{r})]$$

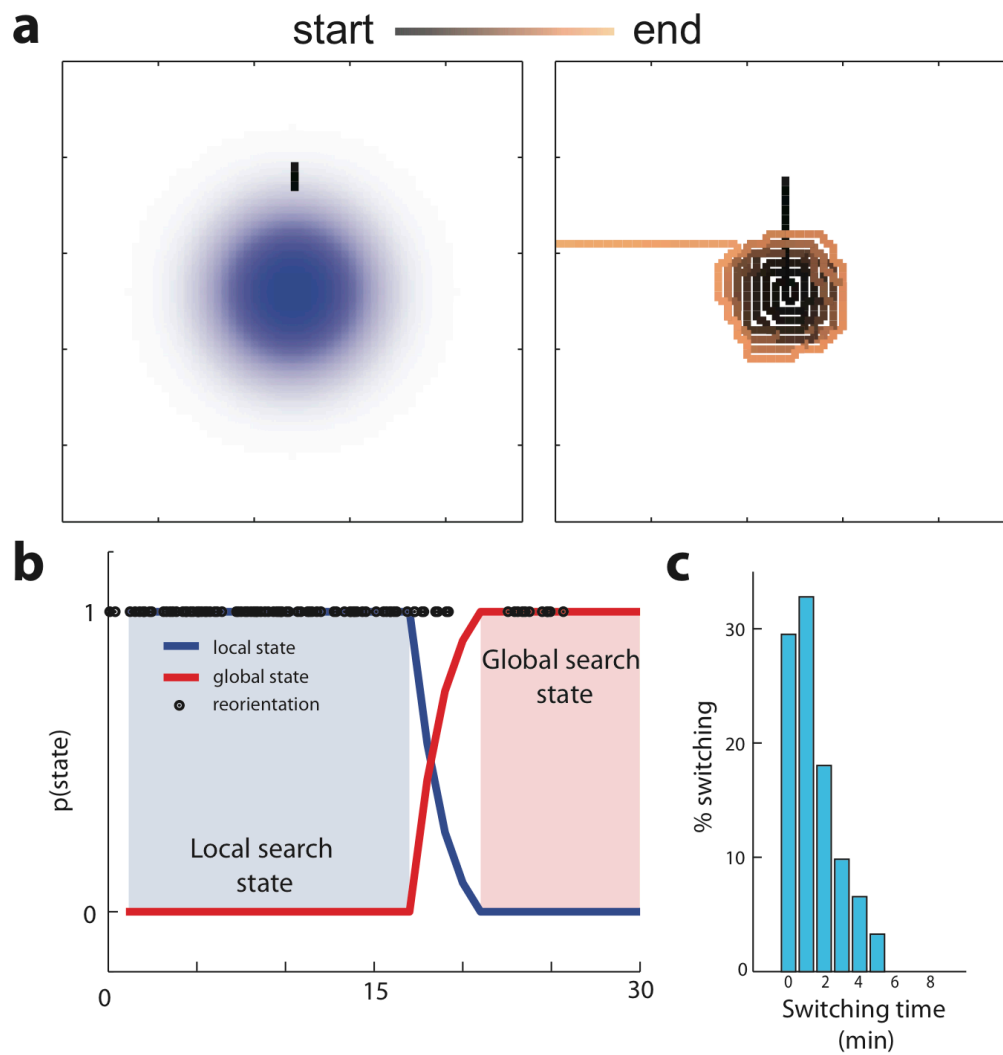


Figure 3.2. Abrupt transitions. a, Initial trajectories of the model head directly towards the peak probability of finding an odor source (*left*). However, after some period of time the model displays an abrupt transition in behavior (*right*). b, Utilizing a Hidden Markov Model to detect transitions in state reveals c, fast switching time across animals.

is updated in a Bayesian manner and can decrease to values below one. Specifically, while the searcher at each step expects to detect a certain number of odorants, none are detected because the source is absent. Therefore, the overall probability that the food source is located within A is updated as follows:

$$p_{t+1}(A|n = 0) = \frac{p_t(A)P(n=0|A)}{P(n=0)},$$

with $p_{t+1}(A) = p_{t+1}(A|n = 0)$. While initially $p_0(A)$ is set very close to 1, this value reaches zero at the transition point between the circular and straight part of the search trajectory in Fig. 3. It is noteworthy that if the probability distribution within A is normalized to one, $p_t(A) = 1$, then the circular trajectory continues indefinitely approximately following an Archimedian spiral (Barbieri et al., 2011). Thus, the transition from local to global search coincides with the searcher decision, at least in the model, that the food is located elsewhere.

The transition between the local and global parts of the search in the model occurs abruptly. Since it is possible that the animal may simply be adapting slowly to a mean gradient, in which case the transition between searches would occur over many minutes. To investigate the sharpness of this transition within the individual trajectories, we applied a hidden Markov model framework (Miller and Katz, 2010; Jones et al., 2007; Seidemann et al., 1996; Abeles et al., 1995; Bishop, 2004) to individual worm trajectories. If segments of single-animal trajectories represent mixtures of states corresponding to local and global parts of the search, then the probability of global search part of the model will increase gradually. The transition

between local and global search states occurs however, on the order of ~ 1 min, which is within the resolution imposed by the frequency of turns (Fig. 3.2c). This is significantly different from a simulated exponentially decaying poisson emitter (Kolmogorov-Smirnov, $p < .05$). Example of this transition for an individual worm trajectory can also be seen in Fig. 3.1a. Thus, the search trajectories both in the experiment and theory exhibit a sharp transition between the local and global parts of the search.

Next we examined whether the infotaxis framework could account for quantitative characteristics of worm search trajectories, such as the probability distribution worm positions at the end of the local search and the cumulative distribution of the durations of local search. The infotaxis model contains three independent parameters: odor diffusivity, the width of the initial prior probability distribution, and how close $p_0(A)$ is set to 1 (see Methods). Fitting these parameters of the infotaxis model, it is possible to account for the experimental distribution of worm positions quantitatively (Fig. 3.3a). Importantly, the same set of the parameters can also account for the cumulative distribution of the durations of local search (Fig. 3b, two-sample Kolmogorov-Smirnov test, $p=0.45$). The conversion between the spatial axis in Fig. 3.3a and the temporal scale in Fig. 3.3b is set by the known value for the worm speed (~ 2 mm/sec), and does not represent an adjustable parameter. Thus, the infotaxis model can quantitatively account for the properties of worm search behavior after removal from food.

The cornerstone feature of infotaxis search trajectories computed in the absence of odorants is that the search reflects the animal's prior beliefs about how the

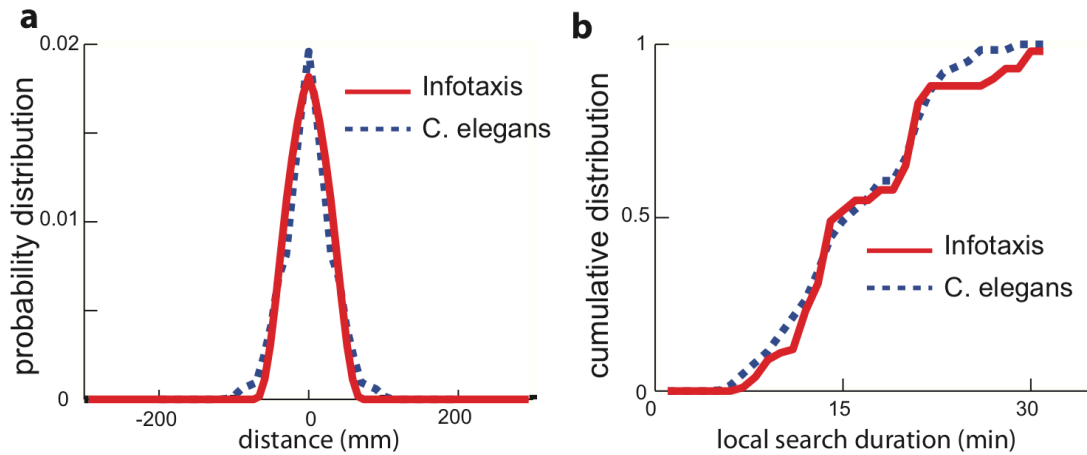


Figure 3.3. Animal behavior matches Infotaxis. a, The distribution of trajectories is matched between the model and behavior. b, This match also reveals a prediction of transition times that is consistent between the individual animals and the algorithm (Kolmogorov-Smirnov, $p=0.45$).

food is likely to be distributed in space. This assumption leads to a prediction that the worm search trajectories should depend only on the relative distribution within the plates with food from which they have been transferred, but not on the overall magnitude of food concentration. This prediction contrasts with predictions based on a more familiar chemotaxis search strategies. The chemotaxis approach would account for the increased number of turns that animals make after they are removed from food as the response to a decrease in odorant/food concentration. Thus, to distinguish chemotaxis versus infotaxis based strategies, we performed the experiment where the worms were acclimated on plates that had the same distribution of food, but at different concentrations. The number of turns made by the worms after removal from plates with food in these two conditions was unchanged (Fig. 3.3c). These experiments thus argue that the worm behavior is more in agreement with the infotaxis model than the chemotaxis model.

The results presented so far argue that the animals' behavior matches that of an optimal maximally informative (Chalasan et al., 2007), or equivalently Bayesian (Najemnik and Geisler, 2005) model. At first glance, these calculations require that the animal maintains, and continuously updates, a "mental map" describing the likelihood to find food throughout the environment. While it seems unlikely that such extensive calculations are performed by *C. elegans*, perhaps the animals use approximations to the full solutions that result in near-optimal performance. The drift diffusion model (Ratcliff and McKoon, 2008) is a classical model of decision-making that allows for plausible neural interpretations of its parameters. Therefore, we sought to determine whether it would be possible to approximate the main features of the maximally

informative solutions, such as the transition time between the local and global parts of the search, using the drift-diffusion model. We found that the quantity $\ln(1 - p_t(A))$, which represents the logarithm of the probability that the food is located elsewhere, follows a drift-diffusion trajectory (Fig. 3.4a). By extracting the mean and variance (drift rate and diffusion, respectively), a drift-diffusion model was able to approximate the transition from local to global search of the full infotaxis model (Fig. 3.4b). It is possible that different parameters of the model use qualitatively different strategies. However, it is possible to rescale the transition times as a function of the width of the prior probability distribution or posterior filter (Fig. 3.4c,d). This suggests that an approximation of the informationally-optimal algorithm is implementable in a simple neural system. Using knowledge of the worm diagram, we suggest a possible neural network model that can implement this strategy (Fig. 3.5). Given that drift-diffusion model conforms well to the standard model of neural spike generation (Tuckwell, 1998), these observations pave the road to neurophysiological studies, perhaps using calcium imaging, of the neural mechanisms responsible for the integration of signals during the local part of the search.

3.4 Discussion

In this work we have shown that exploratory behavior of a small animal, nematode *C. elegans* meets quantitative benchmarks expected for an optimal, maximally informative strategy. This maximally informative strategy involves continuous updates to the likelihood of food sources throughout the environment

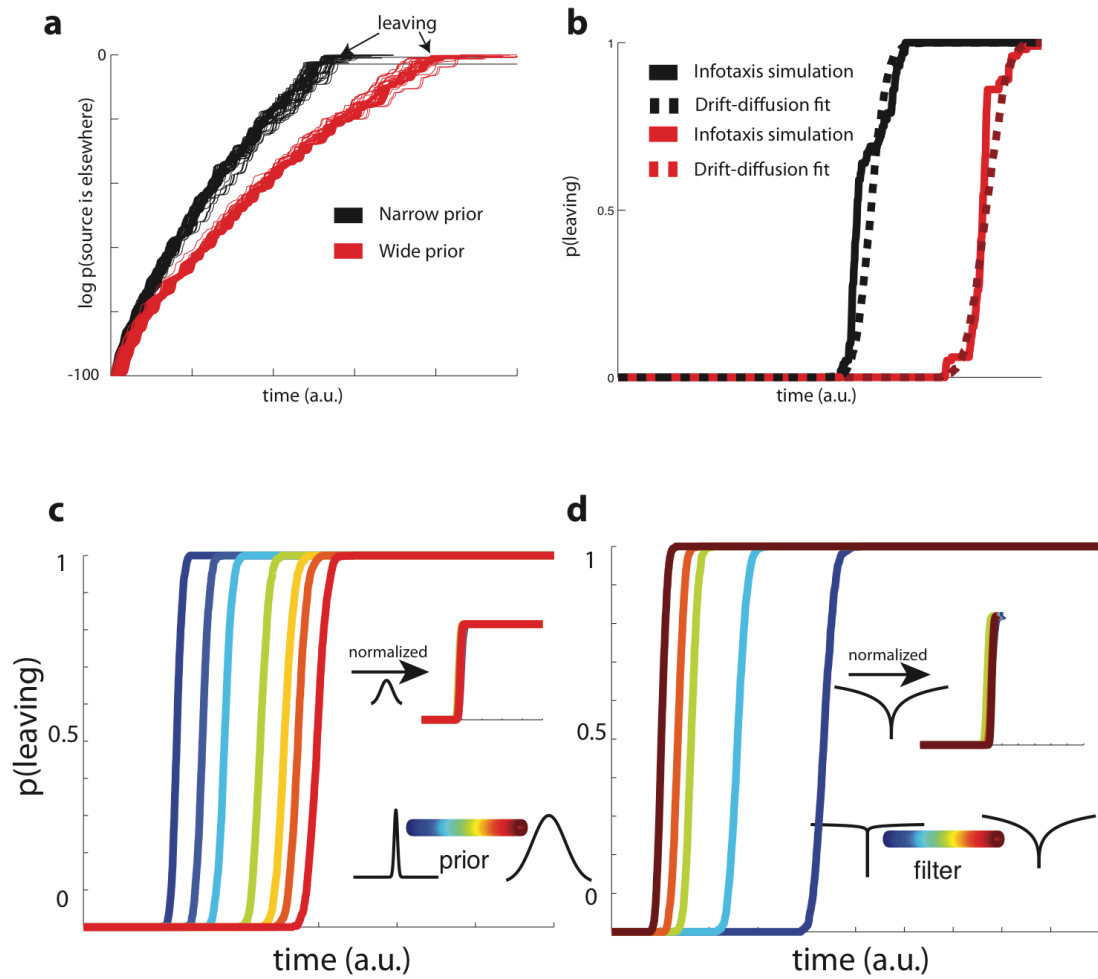


Figure 3.4. Infotaxis is reducible to a drift-diffusion model. a, The log probability that the food is elsewhere resembles a drift-diffusion decision-variable. b, Modeling this evidence accumulation with a boundary at 0 reveals that the model may be represented as an evidence accumulator. The strategies do not fundamentally change as it is possible to normalize by c, prior width or d, filter width.

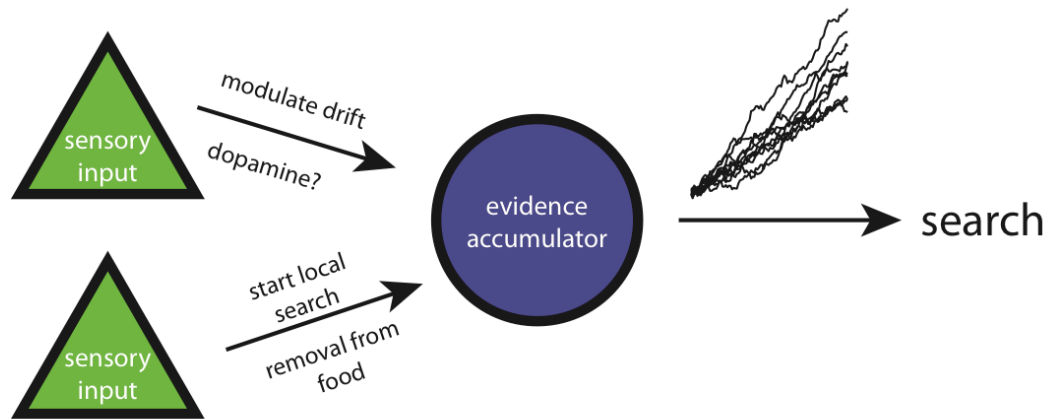


Figure 3.5. Proposed neural model. a, We propose that one could implement this decision-process in a simple neural circuit.

based on incoming sensory inputs. The analysis shows how discrete decisions can naturally arise within continuous behavioral paradigms. Crucially, the animals search trajectories were not affected by the overall concentration within the initial patch of food from the animals were transferred. These observations provide strong evidence that animals' follow an exploratory strategy more complex than a standard chemotaxis approach (Berg and Purcell, 1977).

Given the relative algorithmic complexity of the infotaxis approach, one may wonder how this strategy can be simplified to achieve good performance while minimizing computational cost (Tishby and Polani, 2011). At one extreme one may consider different types of random walk models that may maximize yield under certain distribution of prey. For example, Levy flight random walk have been proposed as good foraging strategies in conditions where food is sparsely distributed, while random diffusive walks are optimal when food is abundant (Viswanathan et al., 2011; Bartumeus et al., 2002; Humphries et al., 2010; Humphries et al., 2012). Both types of strategies do not incorporate recent sensory inputs to guide the behavior. At the other extreme, we have infotaxis model that is presumably integrating information for infinitely long time in the past. These arguments suggest that perhaps neural circuits seek to follow an intermediate strategy that maximizes the gain in performance (Tishby and Polani, 2011) for given investment in algorithmic complexity and neural implements. In particular, the drift-diffusion approximation to infotaxis model suggests that animals might be able approximate the full infotaxis solution just by monitoring one quantity – the probability that the food is located outside of the immediate region under current investigation. This probability follows

a drift-diffusion model (Fig. 3.4) and thus can be easily encoded in the responses of just a single neuron (Tuckwell, 1988). It remains to be investigated how signals from a sensory neuron, such as AWC, would need to be transformed to represent this probability. At the same time, these observations provide a link to previous observations of invertebrate navigational capabilities. Specifically, ants (Wehner and Menzel, 1969) and other invertebrates are known to be able to keep track of their spatial position to return to their nest along the shortest path after foraging search. However, in the case of ants more detailed studies show that the animals use vector summation rather than maintain full spatial maps (Wehner et al., 2006). Our results expand upon these results to show that invertebrates can integrate more abstract quantities than spatial position, although in both cases the transformation from sensory signals to control variables for navigation remain to be investigated.

The mapping between a maximally informative and drift-diffusion models is important, because it relates two powerful frameworks that up to now have been primarily applied in different domains. Most of the work demonstrating the importance of information maximization in the nervous system has pertained to sensory systems, while drift-diffusion models were traditionally used to describe performance in decision-making tasks. A demonstration that the two models can closely approximate each other, provides a principled way to set parameters of drift-diffusion model using information theory and suggests ways for how maximally informative strategies can be implemented in the brain.

The maximally informative framework for describing explorative behavior generates a number of other novel predictions. One prediction is that the duration of

local search should be inversely proportional with the animals' size. Another prediction is that the duration of the local search should be affected by the parameters of the initial distribution of food where animals were grown. This could include both statistical parameters of the distribution, such as its spatial variance, or with improvements in imaging technology, the complete time history of food encounters by individual animals. Testing of these predictions offers hope of building comprehensive models of how past sensory experience influences future behavioral decisions.

Setting parameters of phenomenological drift-diffusion model

3.5 Materials and methods

3.5.1 Quantification of animal behavior

C. elegans in the L4 larval stage were allowed to grow overnight on an agar plate containing a 100 μ l circular patch of the *E. coli* OP50. When it was time for analysis, worms were moved to a clean agar plate, allowed to move at least 3 body lengths away, and finally moved to an agar observation plate. The observation plate contained a 1" border of 200 mM CuSO_4 , which worms will generally not cross. *C. elegans* who have been picked up generally do not like the metal object required to move them and spend roughly two minutes moving forward to escape before starting their area restricted search. Worm movement was recorded for 30 minutes at 3 frames per second, and the first minute is ignored.

3.5.2 Computation of infotaxis trajectories

Infotaxis trajectories were modeled using a 128x128 grid representing position and probability. Although initial descriptions of the model suggested that the reduction in entropy associated with each discrete number of hits should be computed, in our experience any approximation of this results in unwanted behavior. As such, each change in entropy is calculated using the probability to receive a hit or the probability to receive zero hits. Computations are halted upon being within one space of the border or after no movement for 15 time steps. Otherwise, trajectories are computed as in (Vergassola et al., 2007).

3.5.3 Setting parameters of phenomenological drift-diffusion model

During computation of infotaxis trajectories, the value of the decision-variable parameter was saved. The mean and standard deviation of the change in this parameter across all instantiations (initial positions) of the model was computed. Models were run using the equation $dx = A dt + \sigma dW$, where A was the mean (drift) parameter and σ was the standard deviation (diffusion) parameter. Transition times were calculated to occur when the decision variable was greater than or equal to 0.

3.6 Acknowledgments

Chapter 3, in full, is a manuscript in preparation for submission in 2014. Adam J Calhoun, Sreekanth H Chalasani, Tatyana O Sharpee. Emergent decision-making in

information maximizing behavior. In preparation, 2014. The dissertation author was the primary investigator and author of this paper.

3.7 References

- Abeles, M., Bergman, H., Gat, I., Meilijson, I., Seidemann, E., Tishby, N., Vaadia E. Cortical activity flips among quasi-stationary states. *Proc Natl Acad Sci U S A*, 1995. 92(19): p. 8616-20.
- Barbieri, C., Cocco, S., and Monasson, R. On the trajectories and performance of Infotaxis, an information-based greedy search algorithm. *EPL*, 2011. 94:20005.
- Bartumeus, F., Catalan, J., Fulco, U.L., Lyra, M.L., Viswanathan, G.M. Optimizing the encounter rate in biological interactions: Levy versus Brownian strategies. *Phys Rev Lett*, 2002. 88(9): p. 097901.
- Bishop, C.M., *Neural Networks for Pattern Recognition*. 2004, Oxford: Oxford University Press.
- Berg, H.C. and E.M. Purcell, Physics of chemoreception. *Biophys J*, 1977. 20(2): p. 193-219.
- Chalasani, S.H., Chronos, N., Tsunozaki, M., Gray, J.M., Ramot, D., Goodman, M.B., Bargmann, C.I., Dissecting a circuit for olfactory behaviour in *Caenorhabditis elegans*. *Nature*, 2007. 450(7166): p. 63-70.
- Gray, J.M., J.J. Hill, and C.I. Bargmann, A circuit for navigation in *Caenorhabditis elegans*. *Proc Natl Acad Sci U S A*, 2005. 102(9): p. 3184-91.
- Hills, T., P.J. Brockie, and A.V. Maricq, Dopamine and glutamate control area-restricted search behavior in *Caenorhabditis elegans*. *J Neurosci*, 2004. 24(5): p. 1217-25.
- Humphries, N.E., Queiroz, N., Dyer, J.R., Pade, N.G., Musyl, M.K., Schaefer K.M., Fuller, D.W., Brunnschweiler, J.M., Doyle, T.K., Houghton, J.D., Hays, G.C., Jones, C.S., Noble, L.R., Wearmouth, V.J., Southall, E.J., Sims, D.W. Environmental context explains Levy and Brownian movement patterns of marine predators. *Nature*, 2010. 465(7301): p. 1066-9.

- Humphries, N.E., Weimerskirch, H., Queiroz, N., Southall, E.J., Sims, D.W. Foraging success of biological Levy flights recorded in situ. *Proc Natl Acad Sci U S A*, 2012. 109(19): p. 7169-74.
- Jones, L.M., Fontanini, A., Sadacca, B.F., Miller, P., Katz, D.B. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proc Natl Acad Sci U S A*, 2007. 104(47): p. 18772-7.
- Masson, J.-B., Bechet, M.B., Vergassole, M., Chasing information to search in random environments. *J. Phys. A: Math. Theor.*, 2009. 42: p. 434009.
- Miller, P. and D.B. Katz, Stochastic transitions between neural states in taste processing and decision-making. *J Neurosci*, 2010. 30(7): p. 2559-70.
- Najemnik, J. and W.S. Geisler, Optimal eye movement strategies in visual search. *Nature*, 2005. 434(7031): p. 387-91.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922.
doi:10.1162/neco.2008.12-06-420
- Seidemann, E., Meilijson, I., Abeles, M., Bergman, H., Vaadia, E. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci*, 1996. 16(2): p. 752-68.
- Tishby, N. and D. Polani, eds. Information theory of decisions and actions. Perception-action cycle: Models, architectures, and hardware, ed. V. Cutsuridis, A. Hussain, and J.G. Taylor. 2011, Springer: New York. 601-636.
- Tuckwell, H.C., Introduction to Theoretical Neurobiology. 1988: Cambridge University Press.
- Vergassola, M., E. Villermaux, and B.I. Shraiman, 'Infotaxis' as a strategy for searching without gradients. *Nature*, 2007. 445(7126): p. 406-9.
- Viswanathan, G.M. The physics of foraging : an introduction to random searches and biological encounters. 2011, Cambridge ; New York: Cambridge University Press. xiii, 164.
- Wakabayashi, T., I. Kitagawa, and R. Shingai, Neurons regulating the duration of forward locomotion in *Caenorhabditis elegans*. *Neurosci Res*, 2004. 50(1): p. 103-11.

Wehner, R. and R. Menzel, Homing in the ant *Cataglyphis bicolor*. *Science*, 1969. 164(3876): p. 192-4.

Wehner, R., Boyer, M., Loertscher, F., Sommer, S., Menzi, U. Ant navigation: one-way routes rather than maps. *Curr Biol*, 2006. 16(1): p. 75-9.

4. Conclusion

The preceding chapters describe several important advances in how an animal extracts information from the environment over long time windows. Chapter 2 describes how an animal learns about the variability in its environment over long timescales. Chapter 3 describes why an animal would end searching a local area and begin searching a global area. Taken together, these chapters illuminate how an animal extracts and then optimally utilizes information, from neural input to behavioral output.

In Chapter 2, we demonstrate that *C. elegans* uses information about the variability of the environment to actively guide its exploration. We also show that this is detected through a set of specialized sensory neurons that feed into a small circuit. This circuit contains a dopaminergic neuron and is regulated via D1-like dopamine receptors and CREB. The amount of CREB in the cell is one factor in controlling the learning rate, a novel role for CREB. It is interesting to note that the circuit must be a computationally complex one, as removal of any individual neuron is sufficient to block learning. Additionally, the circuit is not a distinct ‘learning module’ but rather a subset of the broader circuit for behavior. The fact that two distinct D1-like receptors appear at different levels of the circuit also suggests possibly unique roles for dopamine at each stage of processing. On the primary sensory neuron, dopamine appears to be acting as a form of gain control.

In Chapter 3, we show that a simple information-maximization strategy recapitulates the transition between local and global search observed in *C. elegans*. This unexpected transition is not encoded into the rules but is rather an emergent

property of the behavior. Interestingly, this model is reducible to a drift-diffusion model of decision-making, something that is neurally-implementable by a small nervous system. The connection between Bayesian inference and drift-diffusion raises the possibility that there may be a rigorous deeper link between these influential theories.

The studies shown here provide several directions for future research. The most vexing question from Chapter 2 remains: what is changing? The system changes through dopamine and CREB, yet it is unclear what is becoming plastic. There may be several potential methods for identifying the mechanism. First, a screen of many mutants may reveal other genes involved in the behavior. Similarly, qPCR of candidate genes could give evidence for change in, say, NMDA-type glutamate receptors. An alternative approach is to take advantage of what we know about the circuitry. Since CREB is acting solely in AIZ and AIB, it may be possible to perform single-cell sequencing only on these cells in order to find the transcripts that vary across conditions.

An additional question concerns the dynamics of the interneurons. How are they responding, and why are so many needed? A tempting direction would be to begin calcium imaging each interneuron individually. However, preliminary investigations reveal difficulties with this approach: some interneurons display different responses depending on the portion of the cell (body, neurite) that is imaged suggesting differing computations. Other neurons display state-dependent responses that are difficult to investigate in the chip. Perhaps a device that can image many neurons simultaneously could trace the flow of information through the circuit.

Alternately, computational modeling may reveal why certain network structures optimally encode and compute information.

The information-maximization strategy also suggests other questions. For instance, how does this extend to other, non-spatial situations? Although the exact form of the equation is specific to spatial situations, the general intuition remains compelling. It would be interesting to see how close human behavior matches infotactic strategies for other tasks. Additionally, for more dynamic tasks information about the current situation may be insufficient for optimal behavior. Thus, incorporation of predictive information into the model would be an exciting possibility.

A final possible direction is the most exciting. Could there be a 1-1 mapping between the infotaxis approximations and the network? In other words, could there be a neuron – perhaps AIB or AIZ – that represents the ‘decision-variable’ of the evidence accumulation model? It need not be a particular neuron, but could be, say, decay of peptide release. The work described here provides new insights into the circuit structure that guides information acquisition and use. It would be fascinating to see each end intimately tied together.