

# Normative approaches to neural coding and behavior

Ann M. Hermundstad<sup>1\*</sup>

<sup>1</sup> Janelia Research Campus, Howard Hughes Medical Institute

\* [hermundstada@janelia.hhmi.org](mailto:hermundstada@janelia.hhmi.org)

## Abstract

These are a brief set of lectures notes for lectures given at the Les Houches summer course in Theoretical Biological Physics in July 2023. In these notes, I provide an introduction to some of the theoretical frameworks that are used to understand how the brain makes sense of incoming signals from the environment to ultimately guide effective behavior. I then discuss how we can apply these frameworks to understand the structure and function of real brains.

Copyright attribution to authors.

This work is a submission to SciPost Physics Lecture Notes.

License information to appear upon publication.

Publication information to appear upon publication.

Received Date

Accepted Date

Published Date

1

## 2 Contents

3	<b>1 Introduction</b>	<b>2</b>
4	1.1 The normative approach	2
5	<b>2 Frameworks that bridge systems</b>	<b>4</b>
6	2.1 Sensory coding	4
7	2.2 Inference	7
8	2.3 Action selection	12
9	2.4 Summary	16
10	<b>3 Closing the loop in a single system</b>	<b>17</b>
11	3.1 Using sensory stimuli to build accurate internal representations	17
12	3.1.1 Maintaining persistent internal representations in the absence of input	18
13	3.1.2 Accurately updating internal representations via internal input	23
14	3.1.3 Reliably tethering internal representations to the external world	24
15	3.2 Using internal representations to guide effective behavior	26
16	3.2.1 Inferring a generative model of behavior	27
17	3.2.2 Learning the parameters of a generative model based on experience	29
18	3.2.3 Building a structured behavioral policy	31
19	3.3 Coupling sensory representations and behavior in closed loop	33
20	3.4 Summary	35
21	<b>4 Outlook</b>	<b>35</b>
22	<b>A Decomposing mutual information</b>	<b>37</b>
23	<b>B Maximum entropy distributions</b>	<b>37</b>

24	<b>C Bayesian ideal observer model in a switching environment</b>	<b>38</b>
25	<b>D Primer on reinforcement learning</b>	<b>39</b>
26	<b>E Derivation of policy gradient algorithm</b>	<b>42</b>
27	<b>F Structuring a motor drive about a goal heading</b>	<b>44</b>
28	<b>References</b>	<b>45</b>
29		
30		

## 31 1 Introduction

32 The brain is responsible for a myriad of different computations that enable animals to behave  
33 effectively in an ever-changing world. The brain must sense signals in the environment, make  
34 inferences and predictions about the causes and consequences of those signals, and must plan  
35 and execute appropriate behavior based on those inferences and predictions.

36 One way that we try to test our understanding of such computations is by building arti-  
37 ficial systems that can mimic the computations and behaviors of real animals. To date, our best  
38 artificial systems pale in comparison to the capabilities of real brains (but large language mod-  
39 els provide an interesting counterpoint); real brains continue to outperform our best artificial  
40 systems. And yet, real brains have many limitations that artificial systems do not—limitations  
41 in speed and precision, and in energy use and storage capacity, for example. In order to over-  
42 come these limitations, it is thought that brains have evolved smart strategies to exploit the  
43 lawfulness of the world in order to achieve “good enough” performance. This can happen  
44 over many difference timescales: over the course of evolution, development, learning, and  
45 adaptation. This idea necessarily implies that there are better and worse solutions for exploit-  
46 ing the lawfulness of the world to secure a performance advantage. One goal of theoretical  
47 neuroscience is to understand the space of such solutions—how many there might be, their  
48 advantages and disadvantages, and, ultimately, why biology might have persisted with one  
49 solution over another. This involves understanding the possible functions of, and constraints  
50 faced by, different biological systems.

51 In the following notes, I broadly organize these computations beginning with sensory input  
52 and moving toward motor output. In between, the brain has to transform sensory input into  
53 appropriate inferences, decisions, and plans of action. In the first section, I begin by discussing  
54 some theoretical frameworks that are used to understand computation across animals and  
55 brain regions. In the second section, I discuss how we can use these frameworks to try to close  
56 the loop in a single system.

### 57 1.1 The normative approach

58 Before diving in, it’s worth briefly discussing the philosophical and methodological approach  
59 that we will take throughout these discussions. Much of theoretical and computational neu-  
60 roscience concerns itself with how the brain works: if I can give you a mathematical model  
61 that accurately reproduces an experimental finding and makes testable predictions for new  
62 experiments, I have understood something about how a certain biological process might un-  
63 fold. A smaller swath of theoretical neuroscience concerns itself with understanding not just  
64 *how* a process might unfold, by *why* it unfolds in this particular way. Why, of all of the possible  
65 strategies that the brain could have adopted to solve a particular problem, is *this* the strategy

66 it adopted? There are many possible answers to this question, some of which could be ‘this is  
67 just what evolution stumbled upon’. But in many cases, the putative answer to this question  
68 helps us understand how a given strategy might be advantageous with respect to performing  
69 a particular function, in a particular setting, and subject to a particular set of biological or  
70 physical constraints.

71 To give a more concrete example: consider your own visual perception of the world. This  
72 perception is something that can be measured and tested by asking you to discriminate dif-  
73 ferent visual stimuli. By designing those stimuli in a way that isolates different statistical  
74 properties of the world, it is possible to examine which properties of the visual world are easy  
75 for you to see, and which are more difficult. This is the purview of the field of psychophysics.

76 For example, if you are presented with visual patterns hidden in a noisy background, you  
77 could be asked to locate the pattern in the noise. The more accurately you can do this, the  
78 better you are able to discriminate these patterns from random noise. It is possible to design  
79 these visual patterns with certain statistical structure by enforcing certain multi-point corre-  
80 lations in light intensity between nearby pixels [1, 2]. Now, if we pick a particular order of  
81 correlation—say, 4th order—it turns out that people can easily discriminate some 4th order  
82 patterns, but struggle to discriminate others, even though they have the same amount of sta-  
83 tistical structure [3]. This is an empirical observation, and we could use this to build a model  
84 that can reproduce this finding. But *why* it is that we should be better able to see some patterns  
85 over others? Why those particular patterns?

86 The answer is that the patterns that we can easily discriminate are those patterns that are  
87 most informative about our natural visual world [3, 4]. Were our visual world organized in  
88 some different way, this would imply that we would be good at seeing different sets of patterns.  
89 Just by knowing the statistical structure of the visual world, we can predict how sensitive  
90 a person will be to different types of visual patterns in the world, and we can do this in a  
91 parameter-free manner. This same idea—that we should be most sensitive to sensory signals  
92 that are most informative about our surroundings—has been used to explain a wide range of  
93 neural and behavioral responses across different species and different sensory systems.

94 Implicit in this statement, however, is the fact that our sensory systems are constrained:  
95 we have to prioritize some sensory signals, some visual patterns, over others. If we had infinite  
96 resources, we would not need to prioritize—we would in principle be able to fully discriminate  
97 any visual pattern from any other. So a more precise version of the previous statement is  
98 that our sensory systems should use minimal resources to maximize the information that they  
99 convey about the natural world in which we operate.

100 This is an example of a ‘normative’ statement: it postulates a goal or a function that a  
101 system is trying to achieve (potentially at the expense of other goals). The mere fact that we  
102 can propose and test normative statements reflects that fact that the world around us obeys  
103 predictable laws that can be exploited by biological systems, and that they systems themselves  
104 are limited in their ability to exploit these laws. If biological systems had infinite memory of all  
105 past states of the world and could perfectly predict and access all future states, there would be  
106 no need to prioritize some goals over others. Similarly, if the world were completely random  
107 and unstructured, there would be nothing to prioritize. In this way, normative theories relate  
108 the lawfulness of the world to the performance advantage that can be achieved by exploiting  
109 that lawfulness. In formulating a normative theory, we postulate (i) a function to be performed,  
110 (ii) a context in which that function will be performed, and (iii) a set of constraints on the  
111 system that performs the given function in the given context. We then derive the optimal  
112 solution for achieving this particular function in this context and subject to these constraints.  
113 In the example above about our own visual sensitivity to different patterns, (i) the function  
114 was to maximize information about incoming visual signals, (ii) the context was the natural  
115 visual world, and (iii) the constraints were limitations in bandwidth (i.e., how many signals

116 could be reliably transmitted downstream).

117 In what follows, we will use this normative approach to consider how sensory systems  
118 might optimally encode and infer features of the external environment, and how downstream  
119 brain regions can in turn guide optimal behavior based on these inferences (Figure 1).

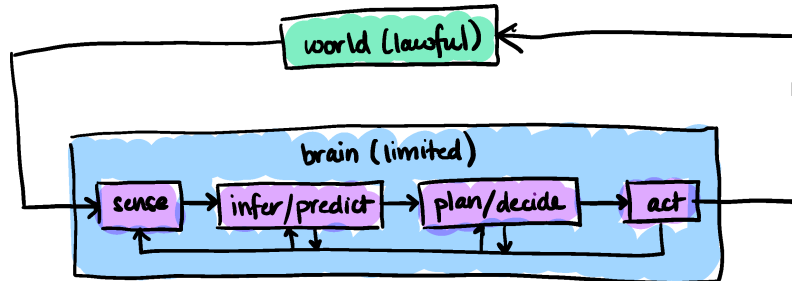


Figure 1: The brain has to use its limited resources to make sense of signals in the environment and guide appropriate actions.

## 120 2 Frameworks that bridge systems

121 In this first section, we will briefly discuss three different normative theoretical frameworks—efficient  
122 coding (Section 2.1), Bayesian inference (Section 2.2), and reinforcement learning (Section  
123 2.3)—that have been used to understand the computational processes involved in sensory  
124 coding, inference, and action selection.

### 125 2.1 Sensory coding

126 The world around us is full of different signals—patterns of light, sounds, smells—that we use  
127 to make sense of our surroundings. At the very earliest stages of processing, the brain must  
128 encode these external signals in internal ones. These internal signals are then used to build  
129 our perceptions of the world, and make decisions based on them. It is remarkable that our  
130 entire understanding of the world is created internally and is built from signals that are first  
131 transduced in our peripheral nervous system.

132 From the perspective of the brain, external signals in the world are not directly accessible,  
133 and can at best be disambiguated from the responses of neurons. This places limitations on the  
134 accuracy with which the brain can “know” the actual state of the environment. To examine the  
135 process of disambiguation, let’s take a brief detour to consider the following example (taken  
136 from David MacKay’s book [5]), which might already be familiar to many readers: you are  
137 given 12 balls that look identical, but one ball is heavier or lighter than the rest. You are given  
138 one scale that can compare any two sets of balls. How do you determine which ball is the odd  
139 ball, and whether it is heavier or lighter, in as few uses of the scale as possible? If you haven’t  
140 already solved this problem before, it’s worth sketching out your solution before moving on.

141 Stated another way, the goal is to design a set of measurements to disambiguate a set of  
142 hypotheses about the world as efficiently as possible (Figure 2). Here, there are 24 hypotheses  
143 in total: there are 12 balls, and each ball could be heavier or lighter, for a total of 24 possibili-  
144 ties. Each measurement corresponds to a usage of the scale, and there are 3 possible outcomes  
145 per measurement: the lefthand side of the scale is heavier, lighter, or equal to the righthand  
146 side. Thus, given 2 measurements, it is possible to disambiguate  $3^2 = 9$  hypotheses; given 3  
147 measurements, it is possible to disambiguate  $3^3 = 27$  hypotheses. From this, we know that it  
148 is possible to disambiguate all 24 of our hypotheses with a total of 3 measurements. The next  
149 step is to design those measurements.

150 The most efficient way to design a measurement is to maximize the entropy over the set of  
 151 possible outcomes for that measurement. Here, by “most efficient”, we mean that the measure-  
 152 ment will reduce our uncertainty as much as possible, given the resolution of the measurement  
 153 device. This is an example of a normative statement: it posits an optimal solution (maximizing  
 154 the entropy over possible measurement outcomes) for a particular problem (disambiguating  
 155 a fixed set of hypotheses as quickly as possible) and subject to a particular set of constraints  
 156 (subject to the resolution of the measurement device).

157 For example, if we begin with 24 hypotheses and have 3 possible outcomes, we want to  
 158 assign  $24/3 = 8$  hypotheses to each outcome. One way to do this would be to divide our 12  
 159 of balls into 3 sets: balls 1-4 (set 1), balls 5-8 (set 2), and balls 9-12 (set 3). We can then  
 160 compare any two sets (e.g. sets 1 and 2). There are then three possible outcomes that evenly  
 161 break apart our hypothesis space:

- 162 1. **Set 1 is heavier.** There are then 8 remaining hypotheses: set 1 has the heavier ball (4  
 163 hypotheses), or set 2 has the lighter ball (4 hypotheses)
- 164 2. **Set 1 is lighter.** There are again 8 remaining hypotheses: set 1 has the lighter ball, or  
 165 set 2 has the heavier ball
- 166 3. **Both sets are equal.** Again, 8 remaining hypotheses: set 3 has the heavier ball, or set 3  
 167 has the lighter ball




Figure 2: Designing measurements to disambiguate hypotheses (adapted from [5]).

168 I will leave it up to the reader to carry this logic through to design the remaining two sets of  
 169 measurements. When applying this measurement technique in succession, we can identify the  
 170 odd ball, and determine whether it is heavier or lighter than the rest, in the fewest possible  
 171 measurements (it’s easy to verify that a different measurement scheme would require more  
 172 than 3 measurements on average). This is precisely because we reduce as much as possible  
 173 our uncertainty about the unknown input—and gain as much information as possible about  
 174 that input—with the outcome of each successive measurement. However, had we been given  
 175 only a single measurement, we would not have been able to fully disambiguate the identity  
 176 and relative weight of the odd ball. For example, if we performed the single measurement

177 shown in Figure 2 and observed that the left side of the scale was heavier, we would at best  
 178 be able to say that one of balls 1-4 was heavier or one of balls 5-8 was lighter. This highlights  
 179 that our ability to disambiguate is limited by the resolution of our measurement device.

180 We can apply a similar perspective to the process of disambiguating signals in the brain.  
 181 For example, consider a neuron that acts as a device to measure incoming light signals. Given  
 182 some knowledge about the distribution of light signals that could be present in the environ-  
 183 ment, one can ask how to design the response function of the neuron to best disambiguate  
 184 these signals subject to the limitations of the device itself. If the neuron can produce a binary  
 185 output (e.g., it can spike or not spike), then this optimal response function should partition the  
 186 incoming distribution of light signals into two equal probability chunks, and assign one output  
 187 to each chunk. This would, for example, lead to a scenario in which the neuron spikes if the  
 188 incoming light signal is greater than the median of the distribution, and is silent if the incoming  
 189 signal is less than the median. By designing a response function that partitions the incoming  
 190 distribution into equal probability chunks, we are guaranteed that the entropy of the neural  
 191 response will be maximized given the distribution of input signals. This is often referred to as  
 192 ‘histogram equalization’, because the histogram of neural responses will be flat. However, it’s  
 193 important to note that there are many ways to partition the distribution into equal probability  
 194 chunks, and thus we often have to invoke other constraints (such as continuity of the response  
 195 function) before comparing to biology.

196 This idea forms the basis of one of the most influential normative frameworks in neuro-  
 197 science: efficient coding. Efficient coding posits that sensory systems maximize the informa-  
 198 tion that they convey to downstream brain regions about incoming sensory signals, and in  
 199 doing so exploit the statistics of the environment in which an organism must function [6].  
 200 This hypothesis was first formulated by Fred Attneave [7] and Horace Barlow [8]. Given an  
 201 input message (here, our stimulus  $S$ ) and an output message (here, our neural response  $R$ ),  
 202 Barlow hypothesized that the goal of the early nervous system is to maximize the information  
 203 that the output message conveys about the input message. The mutual information between  
 204 the input and output is given by:

$$I(R; S) = H(R) - H(R|S) \quad (1)$$

205 where  $H(R)$  is the entropy of the response, and  $H(R|S)$  is the conditional entropy of the re-  
 206 sponse given the stimulus (see Appendix A for derivation). In the limit of low input noise  
 207 (i.e., the stimulus is reliable),  $H(R|S)$  is close to zero, and maximizing mutual information is  
 208 equivalent to maximizing the entropy of the neural response. Barlow used this idea to define  
 209 a quantity called the ‘redundancy’,  $\mathcal{R}$ , to be minimized:

$$\mathcal{R} = (1 - H(R))/C \quad (2)$$

210 where  $C$  is the capacity of the channel along which the messages are sent. It’s worth noting  
 211 that entropy maximization, or information maximization more broadly, does not concern itself  
 212 with how this information will be used downstream. In fact, Attneave discussed this as an one  
 213 of the advantages of the efficient coding framework—that it could be used to understand the  
 214 first stages of sensory processing, without having to know the putative relevance of different  
 215 sensory signals [7].

216 More broadly, this idea can be used to derive the entropy-maximizing set of neural re-  
 217 sponses subject to different constraints on the response distribution (Appendix B). For exam-  
 218 ple, given a constraint on the total number  $N$  of discriminable responses that a neuron can  
 219 produce, the entropy-maximizing response distribution is flat (i.e.,  $P(r) = 1/N$ ; this is our his-  
 220 togram equalization). Given a constraint on the mean firing rate  $\mu$  of the neuron, the response  
 221 distribution is exponential:  $P(r) = \exp(-r/\mu)/\mu$ . And given a constraint on the variance in  
 222 firing rate  $\sigma^2$ , the response distribution is Gaussian:  $P(r) = \exp(-(r - \mu)^2/2\sigma^2)/\sqrt{2\pi\sigma^2}$ .



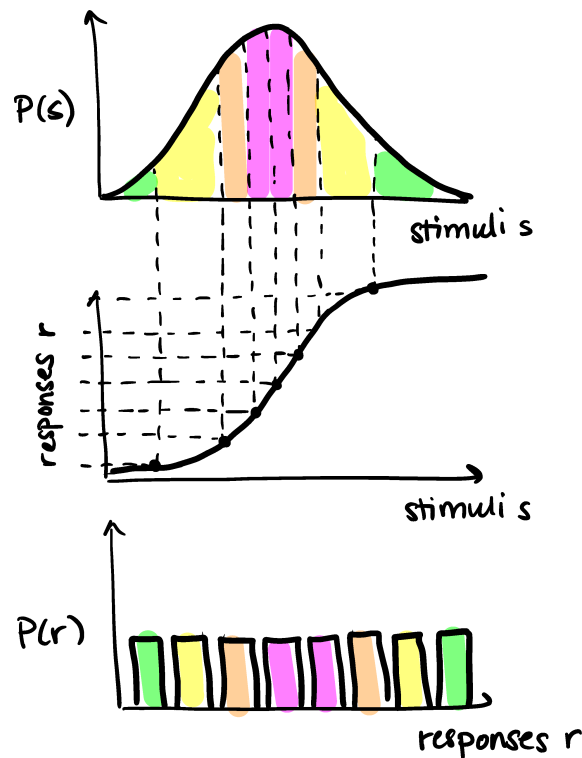


Figure 3: Maximizing entropy about the distribution of incoming stimuli (adapted from [9]).

223 Barlow's hypothesis was first tested experimentally in 1981 by Simon Laughlin [9], who  
 224 showed that he could predict the responses of motion-sensitive neurons in the visual system of  
 225 the blowfly using only the distribution of light signals found in the natural environment (Figure  
 226 3). Since then, there have been numerous studies that have sought to understand different  
 227 aspects of sensory processing in terms of maximizing information about the distribution of  
 228 inputs that a sensory system encounters in its environment, or in terms of minimizing error  
 229 in the reconstruction of those inputs [6]. The optimal sensory coding scheme depends on the  
 230 specific statistics of the environment, the resource limitations and sources of noise within the  
 231 encoding scheme, and the specific objective function used to optimize the encoding scheme.  
 232 As a result, there is no single characterization of an efficient code. For example, in the limit of  
 233 low input noise, when incoming sensory signals are reliable, optimal coding strategies tend to  
 234 decorrelate these signals (consistent with Barlow's redundancy reduction hypothesis). In the  
 235 limit of high input noise, optimal coding strategies tend to average incoming signals in order  
 236 to combat the adversarial effects of noise. These different coding regimes emerge from the  
 237 same underlying framework but under different assumptions, and are predicted at the level of  
 238 single cells [10–13], population codes [14, 15], and behavior [4].

239 A key assumption underlying all of this work is that sensory systems have evolved to ex-  
 240 ploit the particular statistics of the environment in which an organism must function. If that  
 241 environment were to change, we would expect sensory systems to change as well—a point  
 242 that we will examine in the next section.

## 243 2.2 Inference

244 *This section covers material from Młynarski & Hermundstad (2018) and (2021).*

245 In the previous section, we discussed how sensory systems could have evolved to exploit the

246 statistics of the environment. This implies that if the environment were to change, sensory  
 247 systems should change with it (Figure 4). In this way, efficient coding has provided a normative  
 248 perspective on sensory adaptation [16, 17].

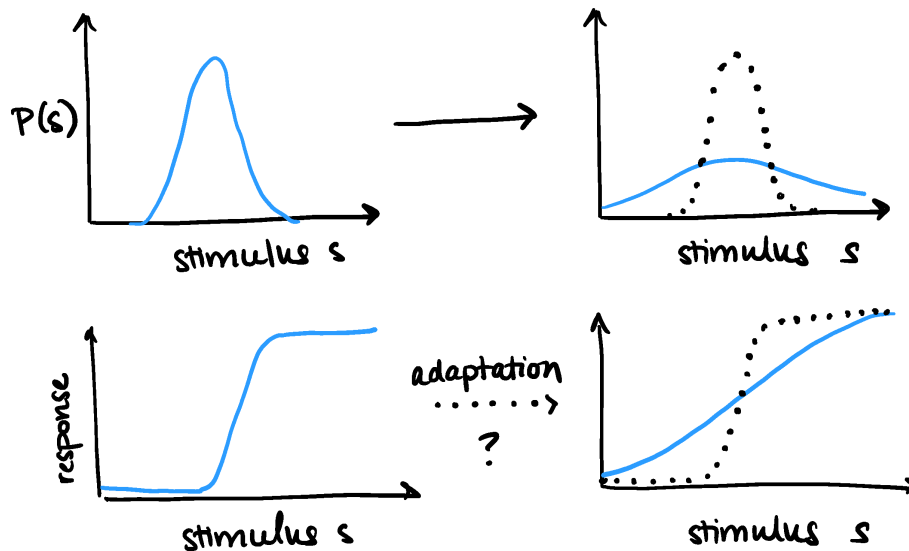


Figure 4: If the statistics of incoming sensory signals change over time, neural responses should adapt to the change.

249 However, the statistics that are most relevant for adapting to a change in the environment  
 250 are not necessarily the same as those that should be optimally encoded in steady state [18].  
 251 Moreover, the different sensory signals carry different relevance for downstream computations,  
 252 which can necessitate different coding strategies [19]. In this section, we treat both of these  
 253 questions.

254 In the previous section, we considered how to best design a neural response function in  
 255 order to efficiently encode signals  $s$  drawn from a distribution  $P(s)$ . We will now assume  
 256 that this distribution can be parameterized by a latent variable  $\theta$ , such that we can specify  
 257  $P(s|\theta)$  (we will also refer to  $\theta$  as the “state” of the environment). For example, for a Gaussian  
 258 stimulus distribution,  $\theta$  can parameterize either the mean or the variance of the distribution. We  
 259 can then compare a neural response function that maximizes the information about incoming  
 260 sensory stimuli (as discussed in the previous section) with a response function that maximizes  
 261 information about the underlying latent parameters of the distribution. In other words, when  
 262 designing the optimal response function, we can concern ourselves with disambiguating the  
 263 current sensory stimulus  $s$  or the underlying latent state  $\theta$ .

264 The parameter  $\theta$  is not observable, and thus its value must be inferred from incoming  
 265 sensory signals. Such a computation falls within the purview of Bayesian inference, which  
 266 specifies the statistically-optimal computation for inferring a latent variable from a set of mea-  
 267 surements. We will consider a simple case described above, in which  $\theta$  specifies either the  
 268 mean or the variance of a Gaussian distribution. We will also assume that  $\theta$  takes on one of  
 269 two values—a high value or a low value. Finally, we will assume that the environment can  
 270 dynamically switch between these two different values with a small but fixed probability  $h$   
 271 per timestep (Figure 5). At each timestep  $t$ , a signal  $s_t \sim p(\cdot|\theta_t)$  is sampled from a distribu-  
 272 tion parameterized by  $\theta_t$ . This setup mimics one that is widely used in sensory neuroscience  
 273 to study the dynamics of adaptation to changes in the mean and variance of sensory stimuli  
 274 [16, 17, 20–23]. The goal of the inference, then, is to use the history of past sensory signals  
 275  $s_{\tau \leq t}$  determine whether  $\theta_t$  is high ( $\theta_t = \theta^H$ ) or low ( $\theta_t = \theta^L$ ) at a particular timestep.

276 To illustrate the inference process, we will derive an “ideal observer” that has certain knowl-






Figure 5: Example nonstationary environment in which the mean (or variance) of a Gaussian stimulus distribution switches between a low and high value over time.

277 edge about the environment, and can use this knowledge to optimally perform the inference.  
 278 In this setting, we will assume that the observer knows (1) that the environment exists in one  
 279 of two states (high or low); (2) the identity of the two states (i.e., the values  $\theta^H$  and  $\theta^L$ ); (3)  
 280 the switching probability  $h$ ; and (4) the values of all other parameters of the stimulus distribu-  
 281 tion (i.e., if  $\theta$  parameterizes the mean of the Gaussian distribution, we assume that the ideal  
 282 observer knows the variance of that distribution; if  $\theta$  parameterizes the variance, we assume  
 283 that the ideal observer knows the mean). In other words, the observer is armed with a great  
 284 deal of prior knowledge; the only thing that the observer does not know is the current value  
 285 of  $\theta_t$ , and the specific signal that will be sampled from  $p(s|\theta_t)$ . To infer  $\theta_t$ , the ideal observer  
 286 can use the history of observed signals  $s_{\tau \leq t}$  to construct the posterior distribution  $P(\theta_t | s_{\tau \leq t})$   
 287 (also called the “posterior belief”, or “belief distribution”):

$$P(\theta_t | s_{\tau \leq t}) = \frac{1}{\Omega} P(s_t | \theta_t) \sum_{\theta_{t-1}} P(\theta_t | \theta_{t-1}) P(\theta_{t-1} | s_{\tau \leq t-1}), \quad (3)$$

288 where

$$P(\theta_t | \theta_{t-1}) = \begin{cases} (1-h) & \theta_t = \theta_{t-1} \\ h & \theta_t \neq \theta_{t-1} \end{cases} \quad (4)$$

289 specifies the probability that the environment switched states at time  $t$ ,  $P(s_t | \theta_t)$  is the stimu-  
 290 lus distribution, and  $\Omega$  is a normalization constant (see Appendix C for a derivation of this  
 291 distribution). For an environment that consists of only two states ( $\theta^L$  and  $\theta^H$ ), the poste-  
 292 rior distribution is defined over two values, and can thus be specified by a single number  
 293  $P_t^L \equiv P(\theta_t = \theta^L | s_{\tau \leq t})$  that specifies the probability that the environment is in the low state  
 294 (and we can use this to compute  $P_t^H = (1 - P_t^L)$ ). We can use this to rewrite the posterior as:

$$P_t^L = \frac{1}{\Omega} P(s_t | \theta_t = \theta^L) [(1-h)P_{t-1}^L + h(1-P_{t-1}^L)] \quad (5)$$

295 From this expression, one can see that the posterior belief that the environment was in the low  
 296 state at time  $t$  depends on the prior belief at time  $t-1$ , weighted by the probability that the  
 297 environment stayed in the low state, and by the likelihood of sampling the observed signal  $s_t$   
 298 in the low state.

299 We can use this belief distribution to construct a point estimate  $\hat{\theta}_t$  of the current state  
 300  $\theta_t$ , for example by computing the mean of the posterior (note that this is the optimal point  
 301 estimate that minimizes the mean-squared error between  $\theta_t$  and  $\hat{\theta}_t$  [24]):

$$\begin{aligned} \hat{\theta}_t &= \langle P(\theta_t | s_{\tau \leq t}) \rangle_{\theta_t} \\ &= \theta^L P_t^L + \theta^H (1 - P_t^L) \end{aligned} \quad (6)$$

302 Standard Bayesian inference assumes that the sensory signal  $s_t$  is directly used to update  
 303 the posterior belief  $P(\theta_t | s_{\tau \leq t})$ . However, as discussed in the previous section, any incoming  
 304 sensory signals must be encoded in neural responses before they are used to perform any  
 305 downstream computations. Because neurons have finite precision and bandwidth constraints,  
 306 there is necessarily loss in this encoding step. As a result, the choice of encoding schemes will  
 307 impact any downstream inferences (Figure 6).

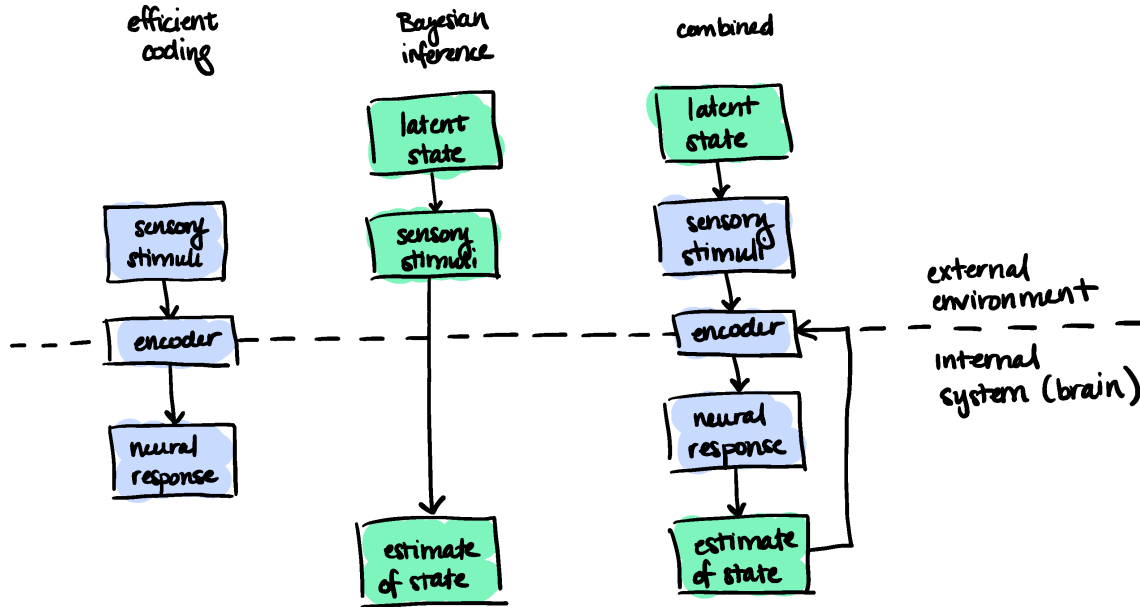


Figure 6: Different architectures of efficient coding and Bayesian inference.

308 To mitigate the negative impact of the encoding, we can design an encoding code scheme  
 309 that preserves information about the incoming stimuli that is relevant for updating the poste-  
 310 rior belief. For example, consider an encoder with a response function of the form:

$$r_t(s_t; k, s_0) = \frac{1}{1 + \exp(-k(s - s_0))} + \eta \quad (7)$$

311 where the parameters  $k$  and  $s_0$  respectively control the slope and offset of this sigmoidal  
 312 function, and  $\eta$  is additive noise (alternatively, we can discretize the output  $r_t$  into a set of  
 313  $N$  discriminable response levels, analogous to the encoding scheme discussed in Section 2.1).  
 314 This response function has finite resolution that it can devote to incoming signals; because of  
 315 the saturating nature of this response function, signals that are sufficiently large or small will  
 316 not be distinguishable from one another. The output of this response function,  $r_t$ , can then be  
 317 decoded to construct an estimate  $\hat{s}_t$  of the incoming sensory signal. For the purposes of this  
 318 discussion, we will assume that we can construct and optimize a simple linear decoder to get  
 319 the estimate  $\hat{s}_t$ . The Bayesian observer must then construct a posterior belief  $P(\theta_t | \hat{s}_{\tau \leq t})$  built  
 320 from the history of past signal *estimates*, rather than directly using the history of true signal  
 321 values.

322 One advantage of choosing a parameterized encoder is that we can directly optimize the  
 323 parameters  $k$  and  $s_0$  of the encoder (as well as the parameters of the linear decoder). For  
 324 the sake of illustration, we will consider two different objective functions for optimizing these  
 325 parameters:

$$\operatorname{argmin}_{k, s_0} \langle (\hat{s} - s)^2 \rangle_{P(s_t | \hat{\theta}_t)} \quad (8)$$

326 and

$$\operatorname{argmin}_{k, s_0} \langle (\hat{\theta}(s) - \hat{\theta}(\hat{s}))^2 \rangle_{P(s_t | \hat{\theta}_t)}. \quad (9)$$

327 The first of these objective functions prioritizes an accurate estimate of incoming sensory sig-  
 328 nals by minimizing the “reconstruction error”; the second of these objectives prioritizes an  
 329 accurate estimate of the latent state of the distribution from which the signals were generated  
 330 by minimizing the “inference error”. In contrast to typical formulations of efficient coding,  
 331 which optimize the encoder based on the true distribution of incoming stimuli  $P(s_t|\theta_t)$ , both  
 332 of these objective functions optimize the encoder based on the *current belief* about the incom-  
 333 ing stimulus distribution,  $P(s_t|\hat{\theta}_t)$  [18, 19]. As a result, any inaccuracies in this belief will  
 334 shape the current encoding scheme, and any inaccuracies in encoding will in turn shape the  
 335 evolution of the posterior belief. This can be viewed as a form of adaptive coding, where the  
 336 output of the encoder is used to update the posterior belief, and the posterior belief is used,  
 337 in turn, to adapt the encoding on the next timestep. Moreover, rather than using the current  
 338 estimate  $\hat{\theta}_t$  to update the encoder, a better strategy is to use a prediction  $\tilde{\theta}_{t+1}$  about the state  
 339 of the environment at the next timestep, and optimize the encoder based on the distribution  
 340  $P(s_t|\tilde{\theta}_{t+1})$ . The point prediction  $\tilde{\theta}_{t+1}$  can be obtained analogously to  $\hat{\theta}_t$ , using the predicted  
 341 posterior:

$$P(\theta_{t+1}|s_{\tau \leq t}) = \sum_{\theta_{t-1}} P(\theta_{t+1}|\theta_t)P(\theta_t|s_{\tau \leq t}). \quad (10)$$

342 As should be expected, the two objective functions in Eqs. 8-9 lead to different optimal param-  
 343 eters depending on the current prediction  $\tilde{\theta}_{t+1}$ . Since this prediction is varying in time, the  
 344 optimal encoding parameters will also vary in time (Figure 7). This can easily be seen through  
 345 numerical optimization. For example, the first objective function leads to an encoder that  
 346 tracks the current prediction of the incoming stimulus distribution: the optimal offset tracks  
 347 the predicted mean of the stimulus distribution,  $s_0^* = \tilde{\mu}_{t+1}$ , and the optimal slope tracks the  
 348 inverse of the predicted standard deviation of the stimulus distribution,  $k^* = 1/\tilde{\sigma}_{t+1}$ . As a  
 349 result, these parameters vary continuously with the current posterior belief (blue curves in  
 350 Figure 7). In the limit that the observer is certain about the current state of the environment  
 351 and maintains a correct estimate of that state, the encoder is aligned with the true incoming  
 352 stimulus distribution, and prioritizes accurate encoding of stimuli from that distribution, anal-  
 353 ogous to classic efficient coding. However, when the state of the environment changes, the  
 354 current encoder is misaligned with the new distribution, and can thus be slow to detect and  
 355 adapt to a change. As a result, when averaged over time, this strategy often suffers from higher  
 356 reconstruction error than a strategy that is optimized for inference. This highlights the fact  
 357 that an encoding that performs optimal reconstruction of sensory signals does not necessarily  
 358 support optimal adaptation to changes in the distribution of those signals.

359 In contrast, the second objective function is optimized for inference and leads to an en-  
 360 coder whose parameters vary discontinuously with the current prediction of the incoming  
 361 stimulus distribution (green curves in Figure 7). The resulting encoder performs a form of  
 362 uncertainty-dependent change detection: when the observer is certain about the current state  
 363 of the environment (and maintains a correct estimate of that state), the encoder shifts its fi-  
 364 nite resolution away from the true stimulus distribution, in anticipation of changes that might  
 365 occur in the future. When the state of the environment changes and the observer becomes  
 366 more uncertain, the encoder takes a form that can best discriminate that change: for mean  
 367 estimation, the optimal encoding function is sharp and centered between the two candidate  
 368 distributions; for variance estimation, the optimal encoding function is shifted toward the tail  
 369 of the high variance distribution (if more certain of the high variance state) or centered about  
 370 the low variance distribution (if more certain of the low variance state). As a result, this en-  
 371 coder is less accurate in steady state, but more sensitive and faster to adapt to changes, and can  
 372 achieve lower average reconstruction error than a strategy that is optimized for reconstruction.

373 As highlighted in the previous discussion, both of these strategies are optimized greedily

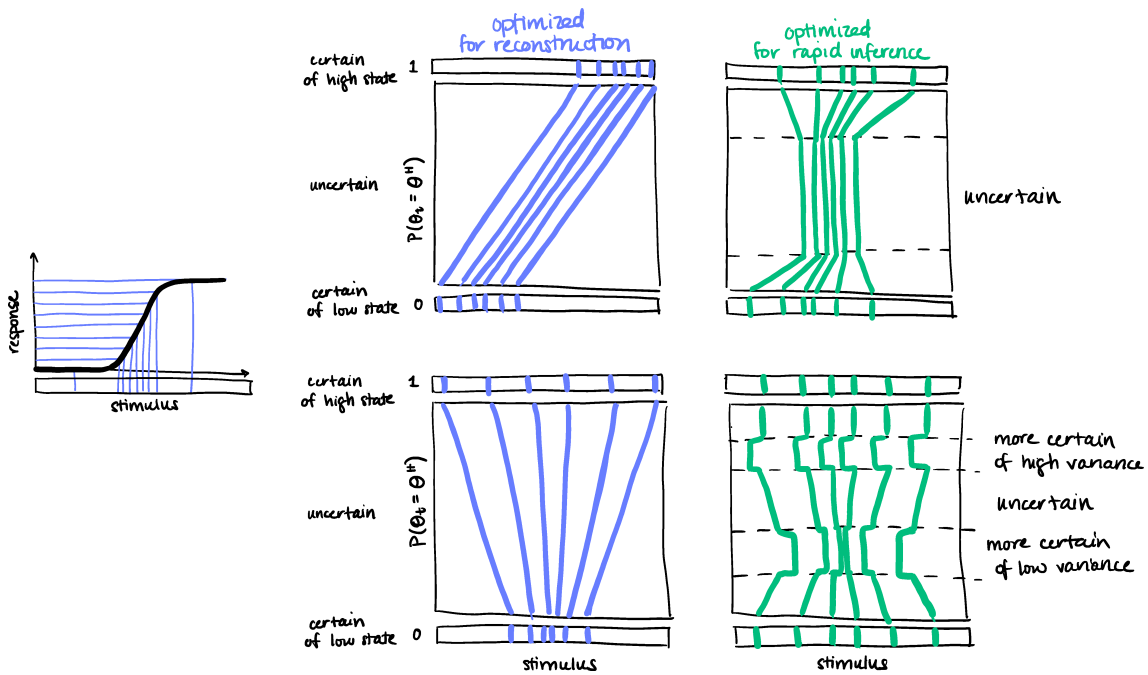


Figure 7: Optimal tuning curves (far left schematic) adapt based on uncertainty, shown for mean adaptation (top row) and variance adaptation (bottom row). Adapted from [18].

374 based on the observer’s current knowledge of the environment, and exhibit different tradeoffs  
 375 in performance. One might therefore wonder whether it is possible to leverage the advantages  
 376 of both strategies. To explore this question, one can form composite codes that balance these  
 377 two objectives:

$$\operatorname{argmin}_{k, s_0} \alpha \langle (\hat{s} - s)^2 \rangle_{P(s_t | \hat{\theta}_t)} + (1 - \alpha) \langle (\hat{\theta}(s) - \hat{\theta}(\hat{s}))^2 \rangle_{P(s_t | \hat{\theta}_t)}, \quad (11)$$

378 where  $\alpha \in [0, 1]$  is a parameter that balances reconstruction and inference errors. For many  
 379 different assumptions about the stimulus distributions and their dynamics, the optimal encoder  
 380 balances these two sources of errors, with an optimal value of  $\alpha$  that is between 0 and 1  
 381 (see [18] for specific examples). This finding suggests that optimally reconstructing signals  
 382 in nonstationary environments requires devoting some bandwidth to detecting changes in the  
 383 underlying distribution of those signals.

384 This also provides a normative perspective on the temporal dynamics of sensory adaptation  
 385 (Figure 8). Sensory codes that are optimized for reconstruction versus inference show differ-  
 386 ent temporal dynamics that mimic those observed in so-called adapting and sensitizing retinal  
 387 ganglion cells in the salamander, mouse, and primate [22, 25, 26]. Moreover, this framework  
 388 predicts that slower environmental dynamics (smaller values of  $h$ ) will lead to slower adapta-  
 389 tion, because the ideal observer requires more stimulus samples to be convinced of a change.  
 390 As a result, the predicted timescale of adaptation scales with the periodicity of changes in the  
 391 environment, consistent with experimental observations in the fly, mouse, rat, and electric  
 392 fish [17, 21, 27, 28].

### 393 2.3 Action selection

394 This section covers material from Ma & Hermundstad (2022)




Figure 8: Optimal neural dynamics exhibit qualitative features of “adapting” (left) and “sensitizing” (right) retinal ganglion cells in response to switches in stimulus variance (adapted from [18]).

395 In the previous sections, we saw how compressed sensory signals (conveyed through a resource-  
 396 constrained encoder) could impact the inference of latent environmental states, and how this  
 397 inference, in turn, could be used to dynamically adapt the encoder over time. We formulated  
 398 this problem in a passive setting, where we did not consider how the inference process im-  
 399 pacted the downstream selection of actions. We will now consider an active setting, where  
 400 inferences guide the selection of actions, which in turn impacts which sensory signals will be  
 401 gathered in the future.

402 We will again consider a scenario in which the environment can take on one of two different  
 403 states,  $\theta_t \in \{\theta_+, \theta_-\}$ , that can change with a small but fixed probability  $h$  per timestep, as given  
 404 in Eq. 4 (note that we have changed our notation from  $R, L$  to  $+, -$  for convenience). However,  
 405 rather than using this state to specify the parameters of a stimulus distribution, we will take this  
 406 state to specify the probability of rewards at two different ports that can be sampled by an agent  
 407 (Figure 9). This is an example of a “nonstationary two-armed bandit” task (more generally  
 408 referred to as a dynamic foraging task) in which an agent is faced with two different levers, or  
 409 “bandits”, that deliver rewards with different probabilities. In a nonstationary setting, these  
 410 probabilities change over time, and thus the agent is best served by using its past observations  
 411 to infer which of the two levers is more rewarding at a particular time. In the simplest version  
 412 of this task, the observations  $o \in \{o_+, o_-\}$  are binary: with each lever pull, the agent either  
 413 receives a reward ( $o = o_+$ ) or receives nothing ( $o = o_-$ ). This type of task has been used  
 414 to study decision making in many different species, including humans [29, 30], non-human  
 415 primates [31, 32], rodents [30, 33–36], and flies [37].

416 At each timestep, the agent has the option of selecting one of two actions,  $a \in \{a_+, a_-\}$ ,  
 417 that correspond to sampling one of the two levers. When the environment is in state  $\theta_+$ , the ‘+’  
 418 lever is more rewarding, and thus the agent should select action  $a_+$ ; when the environment  
 419 is in state  $\theta_-$ , the ‘-’ lever is more rewarding, and the agent should select action  $a_-$ . We  
 420 assume that the more rewarding port delivers reward with probability  $p_{\text{high}}$ , and that the less  
 421 rewarding port delivers reward with probability  $p_{\text{low}}$ . The reward probability can then be  
 422 written as:

$$p(o | \theta, a) = \frac{1 + o (\theta a \Delta p + \Delta \bar{p})}{2} \quad (12)$$

423 where  $\Delta p = p_{\text{high}} - p_{\text{low}}$ ,  $\bar{p} = (p_{\text{high}} + p_{\text{low}})/2$ , and  $\Delta \bar{p} = 2\bar{p} - 1$ , and where we make use of

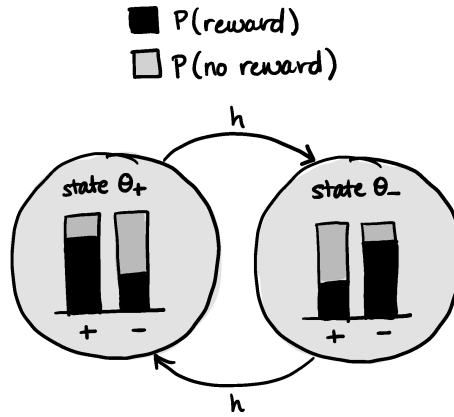


Figure 9: In an example of a nonstationary two-armed bandit task, an agent can sample one of two reward ports ('+' and '-') whose reward probabilities are determined by a latent state  $\theta$  that can switch over time.

424 the following:

$$\begin{aligned} a &\in \{a_-, a_+\} = \{-1, 1\} \\ o &\in \{o_-, o_+\} = \{-1, 1\} \\ \theta &\in \{\theta_-, \theta_+\} = \{-1, 1\} \end{aligned} \quad (13)$$

425 This task is an example of a partially observable Markov decision process, or POMDP. The  
 426 optimal actions in this task can be derived via two separable steps: (1) using Bayesian inference  
 427 to derive the ideal observer that optimally infers a belief  $\mathbf{u}$  about the current environmental  
 428 state, and (2) using reinforcement learning to derive the optimal behavioral policy  $\pi_*(\mathbf{a}|\mathbf{u})$ ,  
 429 which specifies the optimal actions given the current belief [38–40] (see Appendix D for a brief  
 430 primer on reinforcement learning).

431 As in the previous section, we can derive the ideal observer model that optimally infers  
 432 the current environmental state based on its prior knowledge. We assume that the observer  
 433 knows (1) that the environment exists in one of two states (+ or -); (2) the values of  $p_{\text{high}}$  and  
 434  $p_{\text{low}}$ ; and (3) the switching probability  $h$ . The goal of the inference is then to use the history  
 435 of past observations  $\mathbf{o}_{\tau \leq t}$  and actions  $\mathbf{a}_{\tau \leq t}$  to determine the current state  $\theta_t$ . Following the  
 436 derivations in Appendix C, we can write:

$$\begin{aligned} p(\theta_t | \mathbf{a}_{\tau \leq t}, \mathbf{o}_{\tau \leq t}) &= \frac{1}{\Omega} p(o_t | \theta_t, a_t, \mathbf{a}_{\tau < t}, \mathbf{o}_{\tau < t}) p(\theta_t | \mathbf{a}_{\tau < t}, \mathbf{o}_{\tau < t}) \\ &= \frac{1}{\Omega} p(o_t | \theta_t, a_t) \sum_{\theta_{t-1}} p(\theta_t | \theta_{t-1}) p(\theta_{t-1} | \mathbf{a}_{\tau < t}, \mathbf{o}_{\tau < t}) \\ &= \frac{1}{\Omega} \sum_{\theta_{t-1}} \underbrace{p(\theta_t | \theta_{t-1})}_{\text{world dynamics}} \underbrace{p(o_t | \theta_t, a_t)}_{\text{reward delivery}} \underbrace{p(\theta_{t-1} | \mathbf{a}_{\tau \leq t-1}, \mathbf{o}_{\tau \leq t-1})}_{\text{prior at time } t-1} \end{aligned} \quad (14)$$

437 To simplify the final expression of the belief, we define  $\mathbf{u}_t$  to be the difference in belief values  
 438 between the two states:

$$\mathbf{u}_t \equiv P(\theta_t = \theta_+ | \mathbf{o}_{\tau \leq t}, \mathbf{a}_{\tau \leq t}) - P(\theta_t = \theta_- | \mathbf{o}_{\tau \leq t}, \mathbf{a}_{\tau \leq t}). \quad (15)$$

439 Using this, together with the results from Eq. 12 for  $p(o_t | \theta_t, a_t, \mathbf{o}_t)$  and Eq. 4 for  $p(\theta_t | \theta_{t-1})$ ,  
 440 we can write the belief as:

$$\mathbf{u}_t = (1 - 2h) \cdot \frac{a_t o_t \Delta p + (1 + o_t \Delta \bar{p}) \mathbf{u}_{t-1}}{a_t o_t \Delta p \mathbf{u}_{t-1} + (1 + o_t \Delta \bar{p})}. \quad (16)$$



441 Given the belief  $\mathbf{u}$ , we can then determine the optimal value function  $\mathbf{v}_*(\mathbf{u})$  that specifies how  
 442 ‘good’ it is (in terms of future accumulated rewards) to maintain a particular belief  $\mathbf{u}$  (see  
 443 Appendix D for a discussion of value functions). We do this using a technique called value  
 444 iteration [40] that iteratively updates the value function over time based on the outcomes of  
 445 different actions (note that we first must discretize  $\mathbf{u}$  before performing value iteration):

$$\mathbf{v}_t(\mathbf{u}) = \max_{\mathbf{a}} \underbrace{\sum_{\mathbf{u}', \mathbf{o}} p(\mathbf{u}', \mathbf{o} | \mathbf{u}, \mathbf{a}) [\mathbf{o} + (t-1)\mathbf{v}_{t-1}(\mathbf{u}')] }_{q(\mathbf{u}, \mathbf{a})} / t, \quad (17)$$

446 Note that this is modified from standard value iteration to include a running average of ex-  
 447 pected reward [41]; see Eq. D.11 for the standard form. Analogously to the value function  
 448  $\mathbf{v}(\mathbf{u})$ , the action-value function  $q(\mathbf{u}, \mathbf{a})$  quantifies the value of taking a particular action  $\mathbf{a}$   
 449 given a particular belief  $\mathbf{u}$ . These value functions stabilize to their optima relatively quickly;  
 450 we can then use them to specify the optimal policy:

$$\pi_*(\mathbf{a}|\mathbf{u}) = \operatorname{argmax}_{\mathbf{a}} q_*(\mathbf{u}, \mathbf{a}) = \operatorname{sgn} \mathbf{u} \quad (18)$$

451 Note that this corresponds to a purely greedy policy given the current belief:

$$\begin{aligned} \mathbf{a}_{\text{greedy}} &\equiv \operatorname{argmax}_{\mathbf{a}} (r(\mathbf{o} | \mathbf{u}, \mathbf{a})) \\ &= \operatorname{argmax}_{\mathbf{a}} p(\mathbf{o} = \mathbf{o}_+ | \mathbf{u}, \mathbf{a}) \\ &= \operatorname{argmax}_{\mathbf{a}} \frac{1 + \mathbf{a} \Delta p \mathbf{u} + \Delta \bar{p}}{2} \\ &= \operatorname{sgn} \mathbf{u} \end{aligned} \quad (19)$$

452 This illustrates how we can use an optimal belief, derived through Bayesian inference, to de-  
 453 termine the optimal actions conditioned on that belief. In contrast to the previous section,  
 454 here we assume that the agent can precisely measure observations in the environment (a rea-  
 455 sonable assumption if, as in this case, the observations are binary). However, this also assumes  
 456 that agent can precisely track its current belief. Just as we explored how the compression of  
 457 incoming sensory signals impacts inference, so too can we consider how the compression of  
 458 an internal belief can impact the action selection process. However, as soon as we attempt  
 459 to compress the belief  $\mathbf{u}$ , there is no guarantee that the optimal (compressed) policy can be  
 460 derived via the two separable steps of deriving the ideal observer model and deriving the op-  
 461 timal policy conditioned on the observer model. In fact, there are many compressed policies  
 462 that achieve near optimal performance [41] (see [42] for a general treatment of the tradeoff  
 463 between optimal performance and computational complexity).

464 To see this, consider discretizing the belief update in Eq. 16 under the optimal policy in Eq.  
 465 18 (Figure 10). This update can be transformed into a Markov chain that consists of  $\mathbf{n} \in N$   
 466 states corresponding to each of the discretized belief values. These states can be labeled ac-  
 467 cording to the actions that they specify; all states with  $\mathbf{u}(\mathbf{n}) > \mathbf{0}$  are labeled +, and all states  
 468 will  $\mathbf{u}(\mathbf{n}) < \mathbf{0}$  are labeled -. There are  $2N$  transitions between states, corresponding to the  
 469 two observations that can be obtained from taking the action specified by each state; these  
 470 transitions can be determined from Eq. 16. In the limit as  $N \rightarrow \infty$ , this Markov chain will  
 471 approximate the optimal belief update and action selection. However, in the limit that  $N$  is  
 472 finite, the Markov chain can deviate significantly from the optimum, depending on its archi-  
 473 tecture. For  $N$  small, we can directly enumerate all possible Markov chains and compare their  
 474 performance. For  $N$  states, there are at most  $2^N N^{2N}$  possible Markov chains (corresponding  
 475 to 2 labels for each of  $N$  states, times  $N$  target states for each of  $2N$  possible transitions). For  
 476  $N = 5$  states, this amounts to a whopping 312,500,000 possibilities. Luckily, 99.9% of these

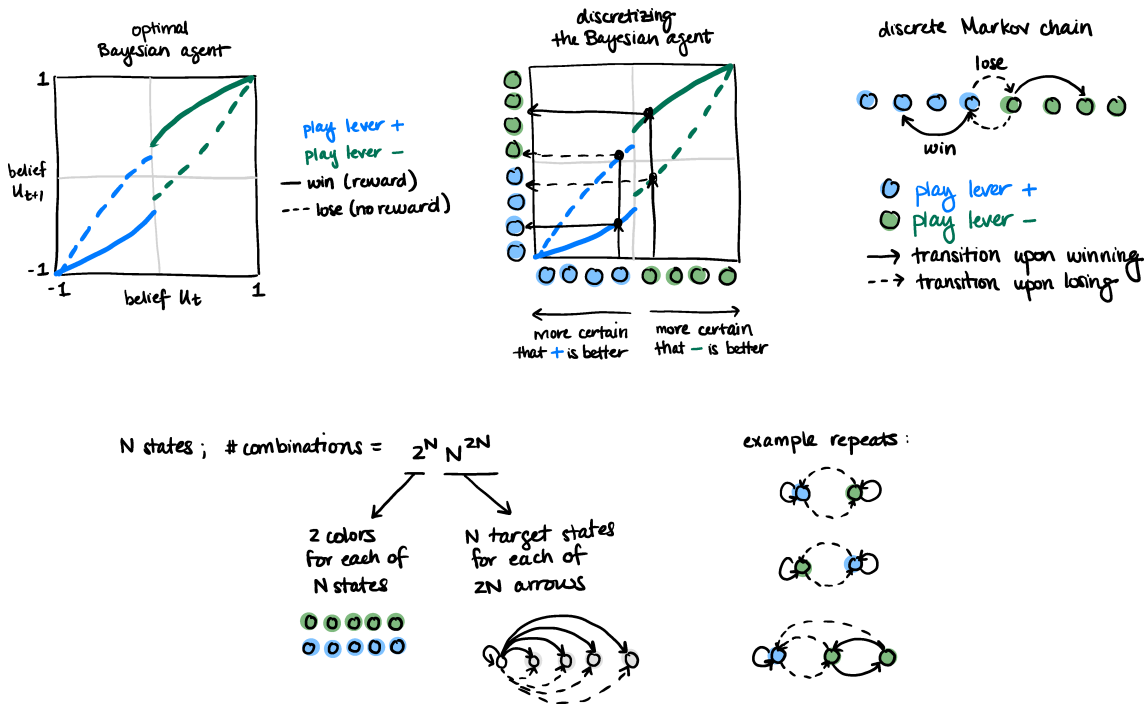


Figure 10: Top row: we can create a compact agent by transforming an optimal Bayesian agent into a discrete Markov chain that specifies actions conditioned on past outcomes. Bottom row: by varying the elements of the Markov chain, there are many ways to build compact agents that achieve good performance (adapted from [41]).

477 are repeats that generate identical behavior under the symmetric task that we've considered;  
 478 a total of **263,428** are unique. This is still a staggering number; again, luckily for us, only a  
 479 fraction of these—on the order of **5,000**—exhibit high performance. Of these, only a small  
 480 fraction exhibit behavior that can be interpreted analogously to the belief update of a Bayesian  
 481 agent [41]. Thus, by considering compact systems that have limited resources for perform-  
 482 ing inferences and guiding actions, we can discover a myriad of solutions that achieve good  
 483 performance.

484 **2.4 Summary**

485 In this section, we have introduced some key normative frameworks—including efficient cod-  
 486 ing (Section 2.1), Bayesian inference (Section 2.2), and reinforcement learning (Section 2.3)—  
 487 that have been used to understand optimal sensory coding, inference, and action selection.  
 488 These frameworks allow us to understand the nature of computations that generalize across  
 489 species, brain regions, and modalities. We have also seen examples of how these frameworks  
 490 inform and constrain one another: for example, in Section 2.2, we saw how different sensory  
 491 coding schemes impact the process of inference, and how inference, in turn, can be used to  
 492 guide adaptive coding schemes; in Section 2.3, we saw how optimal inferences can be used to  
 493 select actions in uncertain and changing environments, and how compact representations of  
 494 those inferences can yield a diversity of different strategies for guiding effective behavior.

495 In the next section, we will use these approaches to understand neural dynamics and be-  
 496 havior in a specific system, and we will follow a set of computations from the encoding of  
 497 sensory stimuli to the selection of actions (and back).

### 498 3 Closing the loop in a single system

499 We now consider a specific system—visually-guided navigation and learning in the fruit fly—  
500 where we can, in principle, follow an entire stream of computations, from the encoding of  
501 sensory stimuli, to their use in building and modifying internal representations of stored vari-  
502 ables, to the selection of actions based on these internal representations. Flies have a rea-  
503 sonably rich behavioral repertoire that allows us to study how behavior changes over time in  
504 response to new experience. Moreover, there is a massive genetic toolkit that enables us to  
505 monitor and manipulate specific cell types during behavior. Finally, with the recent release of  
506 a synaptic-level connectome, it is now possible to relate the physiology of these cell types to  
507 their morphology and connectivity.

508 For the purposes of this discussion, we will consider a class behaviors in which flies asso-  
509 ciate rewards and punishments with stimuli that signal different locations in the environment,  
510 and use these associations to modify their behavior over time, for example to navigate to good  
511 locations or away from bad ones. As an agent in such an environment, this type of navigation  
512 is enabled by knowing (A) where you are, (B) where you want to go, and (C) how to get from  
513 A to B. In what follows, we will consider how these quantities are represented in the brain and  
514 updated based on experience.

#### 515 3.1 Using sensory stimuli to build accurate internal representations

516 *This section covers material from Noorman, Hulse, Romani, Jayaraman, & Hermundstad (2022)*  
517 *and Kim, Hermundstad, Romani, Abbott, & Jayaraman (2019); see also: Fisher, Marquis, D’Alessandro,*  
518 *& Wilson (2019).*

519 We begin by asking how an animal knows where it is—i.e., how the brain keeps track of where  
520 an animal is relative to its surroundings. This is something that many animals do effortlessly.  
521 For example, if I asked you to close your eyes, and with your left index finger, point to the door  
522 that you used to enter the room you’re currently in—you would be able to do this without any  
523 external sensory signals, and without knowing in advance that this is something you’d have  
524 to remember upon entering the room. Our brains use sensory signals to build maps of our  
525 surroundings, and we can use these maps ‘offline’ when we no longer have direct access to  
526 those signals.

527 For several decades now, we have had a theoretical framework that explains how a net-  
528 work of neurons might build and maintain an internal sense of direction [43–51] (see [52–55]  
529 for recent reviews). The theory of ‘attractor networks’ posits that a population of neurons can  
530 encode and update internal representations of a variable (such as which direction you’re fac-  
531 ing) through their recurrent activity. To give some intuition for how this could work, consider  
532 building a neural compass, made from a population of neurons, that keeps track of your cur-  
533 rent direction. There are a few key requirements that we might have for this compass: (1)  
534 we want one single compass needle (no more, no fewer), (2) we want the needle to move  
535 when we turn, and (3) we want the needle to stop when we stop turning. We can build such a  
536 compass with a population of neurons that respond to different orientations; i.e., when you’re  
537 facing north, one group of neurons responds, and when you’re facing east, another group re-  
538 sponds. If we arrange these neurons along a ring based on this tuning, then we can figure  
539 out how to connect these neurons together to meet all of the requirements listed above. To  
540 build a single compass needle, we can connect nearby neurons on the ring with local exci-  
541 tatory connections...but if all neurons locally excite their neighbors, then any activity in one  
542 position along the ring will spread to fill the entire ring, and we will lose our compass nee-  
543 dle. To prevent this from happening, we can have neurons broadly inhibit distant neurons  
544 around the ring; this will keep any activity localized in one region of the ring. With these two

545 ingredients—local excitation and broad inhibition—we can create a single compass needle in  
 546 the form of a localized bump of neural activity (Figure 11).

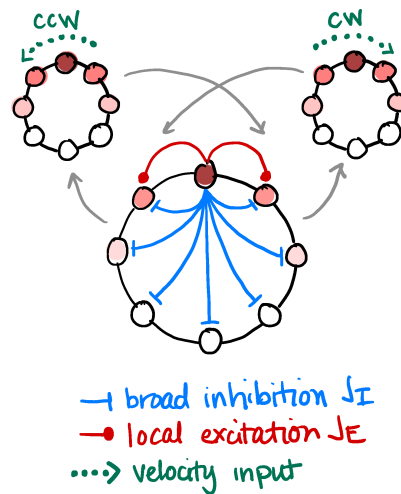


Figure 11: A ring attractor network can maintain a localized bump of neural activity, and can shift the bump by integrating self-motion input.

547 To move this bump of activity, we need to invoke additional mechanisms: we can do this by  
 548 including additional inputs to the ring that are tuned to angular velocity, and that effectively  
 549 ‘push’ the bump around the ring when turning right or left. And to keep the bump still when  
 550 those inputs are removed, we need to ensure that the activity pattern that holds the bump  
 551 is stable, and can persist at the same orientation along the ring without any external inputs.  
 552 Theoretically, this typically requires that we use an infinitely large population of neurons to  
 553 build the compass (see [44, 49, 50, 56] for studies that use large networks to approximate  
 554 the precision of infinite networks, and [57–59] for studies that highlight failure modes of  
 555 small networks). However, biology seems to be able to construct a neural compass from only  
 556 a handful of neurons. The fruit fly, for example, maintains an internal sense of direction  
 557 in a donut-shaped brain structure called the Ellipsoid Body (EB). A population of “compass  
 558 neurons” maintains a persistent bump of activity that encodes the fly’s current heading [60],  
 559 and that can be updated by integrating the fly’s angular velocity [58, 61] (Figure 12). The  
 560 connectivity [62, 63] and dynamics [64, 65] of this network are consistent with theoretical  
 561 accounts of a continuous attractor network, with one notable exception: the fly network is  
 562 composed of a very small number of computational units (on the order of 8 units), rather than  
 563 the extremely large network that we’d expect to need. In what follows, we will explore how  
 564 it is possible to accurately build, maintain, and update an internal sense of direction in such a  
 565 small network.

### 566 3.1.1 Maintaining persistent internal representations in the absence of input

567 To see how we go about building such a network, let’s begin with a population of neurons,  
 568 indexed  $i \in [1, \dots, N]$ , that have neural activities  $r_1, \dots, r_N$ . Let’s consider the case where we  
 569 can describe these activity with a linear system of equations:

$$\dot{\mathbf{r}} = -\mathbf{r} + \mathbf{W}\mathbf{r} + \mathbf{C}_0 \quad (20)$$

570 where  $\mathbf{W}$  is a connectivity matrix that specifies the strength of excitatory or inhibitory con-  
 571 nections between all pairs of neurons, and  $\mathbf{C}_0$  is a constant that specifies feedforward activity  
 572 that is injected into the network. To build some intuition, consider the following connectivity




Figure 12: The fly maintains an internal sense of direction in a neural compass. Left: schematic of the fly brain. Middle: schematic of neurons that innervate a donut-shaped brain region called the ellipsoid body (EB). Right: it is possible to image calcium activity of compass neurons in the EB while the fly is tethered under a microscope and walking on an air-lifted ball. The activity of the compass neurons forms a localized bump that moves around the donut and tracks the movements of the fly relative to the world.

573 matrix:

$$W = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}, \tag{21}$$

574 where  $\lambda_2 < 0$ . This defines a linearly-independent system of equations (adapted from Steven  
575 Strogatz's book [66]):

$$\begin{bmatrix} \dot{r}_1 \\ \dot{r}_2 \end{bmatrix} = \begin{bmatrix} \lambda_1 - 1 & 0 \\ 0 & \lambda_2 - 1 \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \end{bmatrix} + C_0, \tag{22}$$

576 whose solutions are:

$$\begin{aligned} r_1(t) &\propto A_1 + B_1 \exp((\lambda_1 - 1)t) \\ r_2(t) &\propto A_2 + B_2 \exp((\lambda_2 - 1)t) \end{aligned} \tag{23}$$

577 where  $A_1, A_2, B_1, B_2$  are constants. The entries of  $W$  determine the dynamics of this system. If  
578  $\lambda_1 < 1$ , there is one stable fixed point in the network (also called a 'point attractor'); the value  
579 of  $\lambda_1$  (and specifically whether it is less than or greater than  $\lambda_2$ ), determines the relative rate  
580 at which  $r_1$  and  $r_2$  approach the fixed point (Figure 13). Alternatively, if  $\lambda_1 > 1$ , there is a  
581 single saddle point; the dynamics are stable along the  $r_2$  axis, and unstable along the  $r_1$  axis.  
582 In the special case that  $\lambda_1 = 1$ , there is a continuum of stable fixed points along the  $r_1$  axis  
583 (also called a 'line attractor').




Figure 13: Fixed points of the linear system in Eq. 22, with  $\lambda_2 < 0$  (adapted from [66]).

584 In the case where the system of equations is not linearly independent, the same logic  
585 applies to the system of equations. The eigenvectors  $\xi$  of  $W$  specify the axes of flow, and the

586 sign and magnitude of the eigenvalues  $\lambda$  specify the direction and rate of flow, respectively.  
 587 If the system has a single unity eigenvalue and all other eigenvalues less than one, then there  
 588 is a continuum of stable fixed points. This continuum can be used to stably encode a single  
 589 linear variable via appropriate combinations of  $r_1$  and  $r_2$ . This idea was used to propose how  
 590 a network of as few as two neurons, with activities  $r_1$  and  $r_2$ , could stably encode a linear  
 591 variable like the position of an animal's eyes [43].

592 However, this formulation cannot encode a circular variable like orientation. Instead, one  
 593 can build a *ring* of stable fixed points by stitching together multiple line attractors over fixed  
 594 intervals (Figure 14). For this to work, there must be a precise handoff between line attractors,  
 595 which requires a nonlinearity in the network (and some additional fine tuning, which we will  
 596 come to). We can consider a simple nonlinear network of the form:

$$\dot{\mathbf{r}} = -\mathbf{r} + \mathbf{W}\phi(\mathbf{r}) + \mathbf{C}_0 \tag{24}$$

597 where  $\phi(\cdot)$  is a nonlinear function. For illustrative purposes, we will take  $\phi(\cdot)$  to be a threshold  
 598 linear function. This form of nonlinearity ensures that only a subset of neurons in the network  
 599 is active at any given time; importantly, the dynamics of these active neurons are governed  
 600 by a set of *linear* equations via an active submatrix  $\mathbf{W}_{\text{act}}$  of the full connectivity matrix  $\mathbf{W}$ . As  
 601 before, the eigenvalues and eigenvectors of this active submatrix determine the fixed point  
 602 structure of these dynamics. Thus, by appropriately choosing a connectivity matrix  $\mathbf{W}$ , it is  
 603 possible to ensure that all active submatrices of  $\mathbf{W}$  have a single unity eigenvalue, and all  
 604 other eigenvalues less than one. Each active subsystem can thus be constructed to encode a  
 605 line attractor over a finite interval, and these line attractors can then be stitched together to  
 606 form a more complex attractor structure [65].

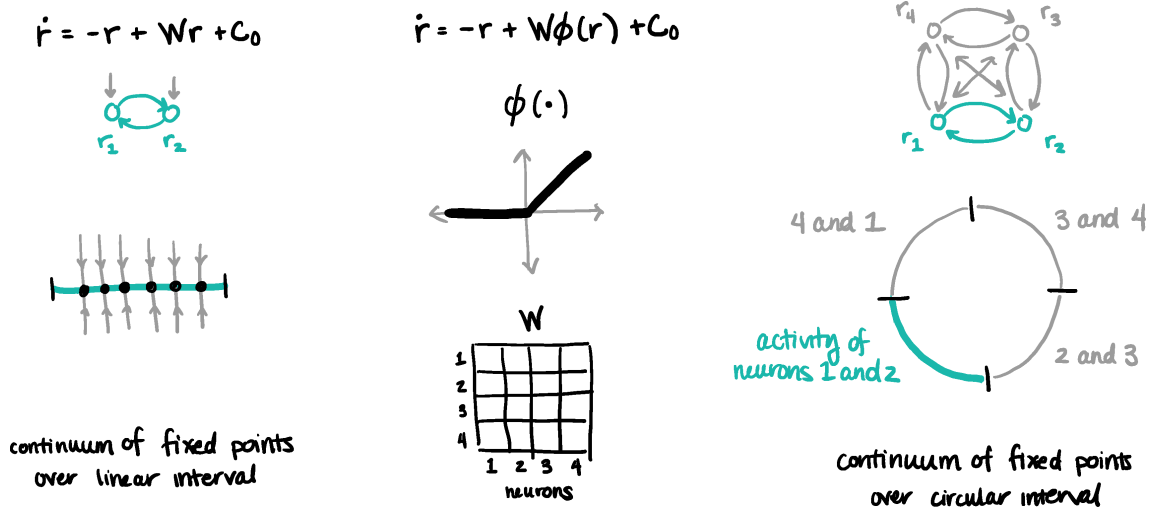


Figure 14: Left: a linear system can encode a continuum of marginally-stable fixed points. Middle, right: in a nonlinear system, a ring attractor can be constructed by stitching together line attractors [65].

607 To construct a ring attractor, we want a rotationally invariant set of line attractors. This  
 608 can be achieved by choosing  $\mathbf{W}$  to be a circulant matrix. A common choice is:

$$\mathbf{W} = \frac{1}{N} (J_I + J_E \cos(\theta_i - \theta_j)), \tag{25}$$

609 where  $J_I$  controls the broad inhibition in the network, and  $J_E$  control local excitation between  
 610 neurons with preferred headings  $\theta_i$  and  $\theta_j$ . With appropriate choices of  $J_E$ , the active sub-  
 611 matrices of  $\mathbf{W}$  are identical, and each has a single unity eigenvalue. For example, given four



612 neurons with activities  $r_1, r_2, r_3,$  and  $r_4$ , a ring attractor can be constructed by stitching together four line attractors spanned by ordered pairs of units: (1)  $r_1$  and  $r_2$ , (2)  $r_2$  and  $r_3$ , (3)  $r_3$  and  $r_4$ , and (4)  $r_4$  and  $r_1$ . In this case, the preferred headings  $\theta_i$  and  $\theta_{i\pm 1}$  are all separated by  $90^\circ$ , and the full connectivity matrix is:

$$W = \frac{J_I}{N} + \frac{J_E}{N} \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{bmatrix}. \quad (26)$$

616 The  $2 \times 2$  active submatrix has a connectivity structure:

$$W_{\text{act}} = \frac{1}{N} \begin{bmatrix} J_I + J_E & J_I \\ J_I & J_I + J_E \end{bmatrix}. \quad (27)$$

617 If we diagonalize this active submatrix, the eigenvectors are  $\xi_1 = [-1, 1]$  and  $\xi_2 = [1, 1]$ , with corresponding eigenvalues  $\lambda_1 = J_E/N$  and  $\lambda_2 = 2J_I/N + J_E/N$  (note that here, we assume  $J_I < 0$ , in which case  $\lambda_1 \geq \lambda_2$ ; more generally, it is possible to derive the maximal values of  $J_I$  for which the network will generate a stable bump of activity [65]). With  $N = 4$  neurons, the leading eigenvalue is 1 if  $J_E = 4$ . This solution guarantees that the bump of activity can persist anywhere along a continuum of orientations between  $\theta_i$  and  $\theta_{i\pm 1}$  (and one can show that this is a solution for all even-sized networks of size  $N \geq 4$ ; see Figure 15 and [65]). The corresponding eigenvector describes a situation in which any increases in the activity of one neuron are offset by decreases in activity of the other, and the resulting population activity pattern is marginally stable. This type of formulation can be used to construct a ring attractor that can encode a continuum of values on a circle with a localized bump of activity. For a network of  $N = 4$ , there is one single setting of the local excitation that can generate a continuum of fixed points with line attractors spanning 2 neurons. For a network size  $N = 6$ , there are 3 ways to generate this continuum from line attractors spanning either 2, 3 or 4 neurons. More generally, for a network of  $N$  neurons, there are  $N - 3$  ways to generate a ring attractor from line attractors spanning  $n = [2, \dots, N - 2]$  neurons. Thus, the larger the network, the easier it is to find a parameter setting that will generate a continuum of fixed points around a ring.




Figure 15: Larger networks have more ways to tune the local excitation in order to generate a ring attractor (adapted from [65]). Circles and squares denote optimal solutions for networks with even and odd numbers of neurons, respectively.

634 This construction assumes that the local excitation,  $J_E$ , can be chosen optimally. If  $J_E$  is not  
 635 chosen optimally, the full connectivity matrix  $W$  will have two different active submatrices that  
 636 differ in size (Figure 16). The larger of these two active submatrices, with a bump spanned  
 637 by  $N_{\text{act}} = n + 1$  active neurons, has a leading eigenvalue greater than one, and thus generates  
 638 an unstable fixed point; the smaller of these active submatrices, with a bump spanned by  
 639  $N_{\text{act}} = n$  active neurons, has a leading eigenvalue less than one, and generates a single stable  
 640 fixed point. The bump of activity will then be pushed away from the unstable fixed point  
 641 and pulled toward the stable fixed point. As it does so, the bump transitions between the  
 642 two different linear regimes that maintain the bump of activity with  $n + 1$  versus  $n$  active  
 643 neurons. The angular span of these two different regimes is closely related to the drift rate  
 644 within each regime—i.e., how quickly the bump is pushed from or pulled toward a fixed point.  
 645 As a result, the dynamics of the bump are governed by three factors: the orientations of the  
 646 stable and unstable fixed points, the rate at which the bump is pulled toward or pushed from  
 647 these fixed points, and the angular span of each stable and unstable regime. In the limit that  
 648  $J_E$  approaches an optimal value, the drift rate in one regime tends to zero as the angular span  
 649 of that regime grows to fill the entire ring. Thus, a ring attractor emerges in the limit that the  
 650 bump drifts infinitely slowly over an increasingly large fraction of the ring.




Figure 16: Left, middle: optimal values of local excitation generate a continuum of marginally-stable fixed points on a ring. Right: non-optimal values of local excitation generate a ring of stable and unstable fixed points.

651 Up until now, we have built intuition about this ring attractor solution by patching to-  
 652 gether linear dynamical systems. An alternative formulation of this solution can be derived  
 653 by computing the “energy landscape” (a Lyapunov function) of the dynamical system (Figure  
 654 17). This energy landscape describes the states to which the system will evolve in the long-  
 655 time limit. In the limit of an infinitely large network with connectivity given by Eq. 25, the  
 656 energy landscape will be flat as a function of orientation. For a finite network with naively-  
 657 chosen  $J_E$ , the energy will be hilly, with discrete basins (corresponding to stable fixed points)  
 658 separated by energy barriers (corresponding to unstable fixed points). The optimal values of  
 659  $J_E$  will flatten this energy landscape. To derive these values, one can first perform a discrete  
 660 Fourier transform on the system of nonlinear equations in Eqs. 24-25 to transform a system  
 661 of  $N$  equations that describe the dynamics of individual neurons into a system of 3 equations  
 662 that describe the dynamics of the amplitude  $a_c$ , width  $w_c$ , and orientation  $\theta_c$  of the compass  
 663 bump (because of the choice of connectivity, only the DC mode and first modes survive the  
 664 Fourier transform, after initial transients). This set of equations can then be used to construct  
 665 the energy of the system. One can then derive the Hessian matrix, which captures the local  
 666 curvature of the landscape. When computed at the orientations of the stable fixed points,

667 the Hessian separates into a block-diagonal matrix, with a single eigenvector aligned with the  
 668 orientation  $\theta_c$ . The corresponding eigenvalue,  $\partial^2 E / \partial \theta_c^2$ , depends only on  $J_E$ ,  $N$ , and  $N_{\text{act}}$ .  
 669 Thus, for a given network size, it is possible to find the values of  $J_E$  that will locally flatten the  
 670 energy as a function of orientation. One can further show that these same values of  $J_E$  agree  
 671 with those derived from the linear systems perspective, and they guarantee that the energy is  
 672 *globally* flat as a function of orientation (see the SI of [65] for more details).




Figure 17: Non-optimal values of local excitation lead to a hilly energy landscape (left); optimal values flatten this landscape (right).

### 673 3.1.2 Accurately updating internal representations via internal input

674 In the previous subsection, we showed how it was possible to maintain a persistent bump of  
 675 activity at any of a continuum of orientations along a ring. For this bump of activity to function  
 676 as a compass needle, it must be able to be updated based on changes in orientation. This can  
 677 be done by injecting self-motion input into the network; the network then has to integrate this  
 678 input to appropriately shift the bump of activity to the correct location along the ring (note  
 679 that this self-motion input is thought to be computed from a combination of motor and optic  
 680 flow inputs [67], which raises an interesting question about whether these inputs are optimally  
 681 combined to estimate self-motion; see [68] for a perspective on this issue). A simple way to  
 682 do this is to add two additional “side rings” to the network that are gated by clockwise or  
 683 counterclockwise angular velocity [69]; these side rings inherit the bump of activity from the  
 684 primary, or “center” ring, and project back to the center ring with shifted connections. Thus,  
 685 clockwise turns will push the bump in one direction around the ring, and counterclockwise  
 686 turns will push the bump in the opposite direction. In the limit that the time constants of these  
 687 side rings are very fast, we can approximate their dynamics with a set of velocity-dependent  
 688 inputs into the network:

$$\dot{r}_j = -r_j + \frac{1}{N} \sum_k (W_{jk}^{\text{sym}} + v_{\text{in}} W_{jk}^{\text{asym}}) \phi(r_k) + C_0, \quad j = 1, \dots, N, \quad (28)$$

689 where  $v_{\text{in}}$  is the input velocity, and where

$$\begin{aligned}
 W_{jk}^{\text{sym}} &= \frac{1}{N} (J_I + J_E \cos(\theta_j - \theta_k)), \\
 W_{jk}^{\text{asym}} &= \frac{1}{N} \sin(\theta_j - \theta_k),
 \end{aligned} \quad (29)$$

690 consistent with Eq. 25 (note that we typically include an integration time constant  $\tau$ , which  
 691 we have taken to 1). If the excitation is chosen optimally, this network can perfectly integrate  
 692 its inputs. However, if the excitation is not optimally chosen, the movement of the bump will  
 693 feel the effects of the stable and unstable fixed points in the network: the bump will move  
 694 faster than the input velocity as it accelerates away from an unstable fixed point, and the  
 695 bump will move slower than the input velocity as it decelerates toward a stable fixed point.  
 696 These dynamics are again determined by the angular orientations of the unstable and stable  
 697 fixed points, the angular span of the linear regimes about each fixed point, and the drift rate at

698 which the bump is pushed away from or pulled toward these fixed points. In the limit of small  
699 velocity inputs, the drift rates and angular spans of each regime will remain unchanged, but  
700 the orientations of the fixed points will shift; the stable fixed points will shift in the direction  
701 of the input velocity, and the unstable fixed points will shift against the direction of the input  
702 velocity. So long as the stable and unstable fixed points remain within their respective regimes,  
703 the bump will evolve toward and persist at the stable fixed point, and the network will fail to  
704 continuously integrate its inputs. However, above a particular threshold velocity, the stable  
705 fixed point will move into the *unstable* regime, and the unstable fixed point will move into the  
706 *stable* regime. As a result, when in the stable regime, the bump will be pulled toward a stable  
707 fixed point, but can never reach it; instead, it will transition into the unstable regime, where  
708 it will be pushed away from an unstable fixed point. This push and pull will cause the bump  
709 to speed up and slow down as it transitions between regimes.

710 In the view of the energy landscape, the drift rates determine the local curvature of the  
711 energy landscape within each linear regime. As the velocity increases, the fixed points will  
712 move, but the local curvature about those fixed points will remain approximately unchanged.  
713 This leads to a tipping of the energy landscape in the direction of the input velocity. If the input  
714 velocity is very small, the degree of tipping will be small, and the bump will still get stuck at  
715 the stable fixed points in the network. As the input velocity increases, the tipping will be more  
716 severe. Again, above a given threshold velocity, the bump will move continuously down the  
717 energy landscape without getting stuck, but will speed up and slow down as it feels the effect  
718 of the curvature of the landscape.

### 719 3.1.3 Reliably tethering internal representations to the external world

720 The previous subsections highlighted how we could build a network that maintains and up-  
721 dates an internal sense of direction, even in the absence of external sensory cues. However, to  
722 function in an external environment, this internal representation must be tethered to cues in  
723 that environment. In the fly compass network, these external sensory cues reach the compass  
724 through another set of “ring neurons” [63]. There are many different classes of ring neurons  
725 that bring in information about polarized light [70–72], wind direction [73], and local visual  
726 features [74–77]. This again raises an interesting question as to whether and how these differ-  
727 ent inputs are combined to (optimally?) infer an estimate of orientation, something we will not  
728 discuss in detail here. For the sake of illustration, we will focus one modality that is encoded  
729 by visual ring neurons that respond to local spatiotemporal features in a scene [74, 76]. The  
730 response properties of these neurons are well captured by Gabor-like spatial filters combined  
731 with biphasic temporal filters, a structure that is thought to efficiently exploit the structure of  
732 spatiotemporal correlations in natural scenes (as discussed in Section 2.1; [78, 79]). At the  
733 simplest level, this means that a given ring neuron will be active whenever a particular feature  
734 in the visual scene is present at a particular location relative to the fly. The presence of multiple  
735 inputs to the compass—here, external sensory inputs and internal self-motion inputs—raises  
736 the question as to how the compass network can reliably keep these sets of inputs in register  
737 to accurately update a single, self-consistent internal representation of heading.

738 This is thought to be achieved through plasticity between ring neurons and compass neu-  
739 rons [80–83]. Each ring neuron makes all-to-all synapses onto compass neurons [63], and  
740 these synapses are plastic and can be modified over time based on experience [82–84]. As a  
741 result, plasticity is thought to “map” the external world onto the compass, and several lines of  
742 ongoing work are focused on how this map remains self-consistent [82–85]. One set of ideas  
743 stems from Kohonen’s ‘self-organizing map’, which is an unsupervised competitive learning  
744 algorithm that tries to iteratively map a high-dimensional input space into a low-dimensional  
745 ‘map space’ [86] (Figure 18). Given an input space spanned by input patterns  $\{p\}$  and a map  
746 space spanned by output patterns  $\{q\}$ , this is achieved through plasticity in weights  $W_{ij}$  that




Figure 18: Illustration of self-organizing map developing over time (adapted from Wikipedia, user:Mclld).

747 link the two:

$$\Delta W_{ij} = \alpha \theta(j)(p_i - W_{ij}) \quad (30)$$

748 For a given input pattern  $\mathbf{p}$ , this rule updates the weights in the local neighborhood  $\theta(j)$   
 749 around the most strongly responding output unit  $q_j$  (also referred to as the ‘best matching  
 750 unit’). The parameter  $\alpha$  is a learning rate that determines how quickly the weights will update  
 751 over time.

752 If we define the input patterns to be the responses of the ring neurons, and the output  
 753 patterns to be the responses of the compass neurons, then the ring attractor enforces a natural  
 754 neighborhood function defined by the bump of compass activity. The rule then becomes:

$$\Delta W_{ij} = \alpha r_j^c (r_i^r - W_{ij}) \quad (31)$$

755 where  $r_i^r$  and  $r_j^c$  define the activity of the ring and compass neurons, respectively. This rule is  
 756 similar to a classic Hebbian learning rule in which neurons that “fire together, wire together”—  
 757 in other words, coactivity between neurons  $i$  and  $j$  will lead to a strengthening of the weights  
 758 that link them.

759 In the case of the fly’s neural compass, the ring neurons are thought to be inhibitory [87,  
 760 88], and so the appropriate learning rule is “anti-Hebbian” (i.e., coactivity between neurons  $i$   
 761 and  $j$  will weaken, rather than strengthen, the weights that link them). Moreover, evidence  
 762 suggests that the learning rate  $\alpha$  scales with the fly’s angular velocity  $v$  [82, 84], and thus has  
 763 a form similar to:

$$\Delta W_{ij} = -v^2 r_j^c (r_i^r - W_{ij}) \quad (32)$$

764 This velocity-dependent learning rate ensures that the map is only updated when the fly turns,  
 765 and does not update when the fly persists at the same orientation for long periods of time.

766 As the fly explores a new visual scene, this plasticity rule will generate a self-consistent  
 767 mapping of the visual world onto the compass, such that turns that drive the bump of activity  
 768 are matched by weak inhibition from ring neurons that respond to the corresponding features  
 769 in the visual scene (Figure 19). For example, consider a visual scene that consists of a single  
 770 visual feature, like a vertical bar. And consider that a given ring neuron  $r_i^r$ , which responds  
 771 to features at an orientation  $\phi_i$ , is co-active with a given compass neuron  $r_j^c$ , which responds  
 772 when the bump is at an orientation  $\theta_j$ . If the fly turns, the visual scene will change by an  
 773 orientation  $\Delta\phi$ , such that a ring neuron  $i + \Delta\phi$  is now active. At the same time, self motion  
 774 inputs will drive the bump to change orientations by some angle  $\Delta\theta$ , such that a compass

775 neuron  $j + \Delta\theta$  is now active. The plasticity rule above will ensure that ring neuron  $i + \Delta\phi$   
 776 will have the weakest weight, and thus the smallest inhibition, onto compass neuron  $j + \Delta\theta$ .  
 777 As a result, the compass bump will move in lock step with the movements of the fly and the  
 778 movements of the visual scene.




Figure 19: Plasticity creates a self-consistent mapping between visual and self-motion inputs that drive the compass (adapted from [82]; see also [83]). Left: visual ring neurons have receptive fields that tile visual space. These neurons make all-to-all inhibitory synapses onto compass neurons that maintain the compass bump of activity. Middle, right: as the fly turns, the compass bump moves around the EB, exciting a different compass neuron (upper row). At the same time, the visual scene moves relative to the fly, and excites a different ring neuron (middle row). The coactivity between ring and compass neurons weakens the synapses between them. Over time, this creates a self-consistent mapping between the visual and self-motion inputs that drive the compass.

779 By this same logic, one can easily construct a scenario that might alter this mapping. For  
 780 example, consider a visual scene with a two-fold symmetry, such as two identical bars separated by  $180^\circ$ . This scene will induce identical responses in two ring neurons that respond to orientations separated by  $180^\circ$ . However, the plasticity rule introduced above will eventually  
 782 weaken synapses at only one location in the compass network, leading the bump to favor one  
 783 of the two orientations that are associated with this symmetry. I will leave it to the reader to  
 784 speculate as to how this might impact the dynamics of the compass bump; we will return to  
 785 this idea in Section 3.3.  
 786

### 787 3.2 Using internal representations to guide effective behavior

788 *This section covers material from Dan, Hulse, Kappagantula, Jayaraman, & Hermundstad (2021).*

789 We can now consider how this internal sense of direction is used to guide the fly's behavior.  
 790 For this, we will start by analyzing the behavior itself, and those changes in behavior that  
 791 arise during learning. We will then explore how these changes in behavior could be mediated  
 792 by internal representations in the brain, such as the sense of direction that we have been  
 793 discussing.

794 For this, we will consider a variant of a older learning paradigm [89] in which flies navigate  
 795 in a one-dimensional virtual environment (Figure 20; physicist readers might find it interest-  
 796 ing to know that these experiments were pioneered by Martin Heisenberg, the son Werner



797 Heisenberg, prompting some to ask whether one can formulate an analogous ‘fliesenberg un-  
 798 certainty principle’ [90]). This environment consists of an LED screen that spans 330° and  
 799 displays different visual patterns [91, 92]; this screen surrounds the fly, which is tethered in  
 800 place but able to move its wings. The orientation of the patterns on the screen is then coupled  
 801 to the fly’s movements; when the fly tries to turn left, the patterns on the screen are rotated  
 802 to the right, and vice versa. In this way, the fly has closed-loop control of the movements of  
 803 the patterns on the screen. By tethering the fly, it is then possible to image from the brain  
 804 while the fly is navigating. By further pairing orientations of the visual scene with rewards  
 805 and punishments, it is possible to study how the fly’s neural activity and behavior change in  
 806 response to different types of feedback.




Figure 20: Flies can learn to avoid punishment associated with different visual patterns (adapted from [93]). Left: a visual scene has repeating sets of visual patterns; one set of patterns is paired with an aversive heat punishment. Right: in a virtual reality setup, the fly orients relative to the visual scene, and is pinged in the abdomen with an infrared laser whenever it orients toward one of the punished visual patterns.

### 807 3.2.1 Inferring a generative model of behavior

808 Flies exhibit different modes of patterned behavior. In flight, they exhibit periods of straight  
 809 flight, or ‘fixations’, punctuated by abrupt turns, or body ‘saccades’ [91, 93–95] (Figure 21).  
 810 We can segment the behavior into these two different modes, measure how the properties  
 811 of these modes vary over time and across flies, and use this to infer a generative model of  
 812 behavior. This generative model can be phrased in terms of a behavioral policy  $\pi(\mathbf{a}|\boldsymbol{\theta})$  that  
 813 specifies the probability of taking an action  $\mathbf{a}$  given a particular state  $\boldsymbol{\theta}$  of the environment [40]  
 814 (analogously to our discussion of behavioral policies in Section 2.3).

815 We begin by assuming that behavior can be structured as a sequence of saccades and fixa-  
 816 tions, such that the termination of one mode initiates the other mode, and vice versa. In the  
 817 language of reinforcement learning, each action then comes in two different types:  $\mathbf{a} \in \{F, S\}$ .  
 818 Each mode can be specified by an angular velocity,  $\boldsymbol{\omega}$ , and a duration conditioned on that  
 819 angular velocity,  $\Delta t|\boldsymbol{\omega}$ .

820 Fixations tend to have near-zero angular velocity, so we can approximate this as a delta  
 821 function:  $P(\boldsymbol{\omega}) = \delta(\boldsymbol{\omega})$ . The distribution of fixation durations—whether measured across  
 822 time or flies—is well fit by an inverse Gaussian distribution:

$$P(\Delta t; \mu, \lambda) = \sqrt{\frac{\lambda}{2\pi\Delta t^3}} \exp\left(-\frac{\lambda(\Delta t - \mu)^2}{2\mu^2\Delta t}\right) \quad (33)$$

823 This distribution is appealing because we can specify a process that generates it. Consider the  
 824 following stochastic process:

$$X_t \sim vt + \sigma W_t \quad (34)$$

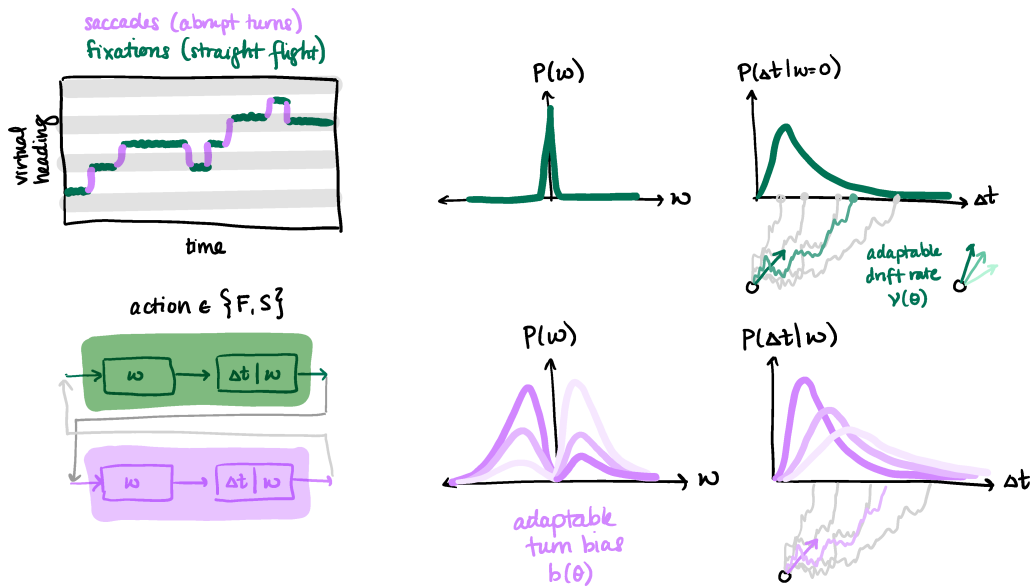


Figure 21: Left: flies fixate (F) and saccade (S). These flight patterns vary in angular velocity  $\omega$  and duration  $\Delta t$ . Right: we can use the statistics of these flight patterns to infer a generative model of behavior (adapted from [93]).

825 where  $W_t$  is standard Brownian motion, and  $\nu$  is a drift term (this process is often called a  
 826 'drift diffusion' process). The first passage time of this process, given a fixed threshold  $\rho > 0$ ,  
 827 follows an inverse Gaussian distribution, with  $\mu = \rho/\nu$  and  $\lambda = (\rho/\sigma)^2$  [96–98]. Thus, we  
 828 can capture a single fixational event by sampling from this process; the fixation terminates  
 829 when the process crosses a threshold  $\rho$ .

830 This process requires specifying three parameters: the threshold  $\rho$ , the drift rate  $\nu$ , and  
 831 the diffusion spread  $\sigma$ . In principle, any or all of these parameters could vary over time based  
 832 on experience, and the variation in each of these parameters would manifest in predictable  
 833 changes in the mean and variance of the inverse Gaussian distribution. To see this, note that  
 834 the mean  $\text{mean} = \mu$  and variance  $\text{var} = \mu^3/\lambda$  are related by:  $\log(\text{var}) = 3 \log(\text{mean}) - \lambda$ . If,  
 835 for example, the variability in behavior is explained by changes in the drift rate  $\nu$  alone (with  
 836 fixed threshold  $\rho$  and spread  $\sigma$ ), we would expect to see the relationship between the variance  
 837 and mean of fixation durations to follow a line with a slope of 3 and offset of  $-\lambda = -\rho^2/\sigma^2$ . By  
 838 fitting this distribution across time and across flies (or, alternatively, by measuring the mean  
 839 and variance empirically), we can show that the behavior is indeed most consistent with a  
 840 model in which the drift rate is changing over time. Thus, we can model fixations as events  
 841 with near-zero angular velocity, and with durations generated by a drift diffusion process with  
 842 an adaptable drift rate.

843 Saccades follow a similar structure, but on a different time scale. In contrast to fixations,  
 844 these events are thought to be ballistic, such that they cannot be interrupted or modified dur-  
 845 ing their execution. Using the same types of analyses, we can show that the angular velocity of  
 846 saccades is well fit with two log-normal distributions (corresponding to clockwise and counter-  
 847 clockwise saccades), but with an adaptable bias that captures the tendency to initiate saccades  
 848 in one direction versus another. We find that the duration of saccades, analogously to fixations,  
 849 also follows an inverse Gaussian distribution. However, rather than having a flexible drift rate,  
 850 the distribution of saccades is consistent with an angular-velocity-dependent drift rate. Thus,  
 851 we can model saccades as events with nonzero angular velocity drawn from a lognormal distri-  
 852 bution with an adaptable bias, and with a duration generated by a drift diffusion process with  
 853 an angular-velocity-dependent drift rate. The similarity in the statistical structure between

854 fixations and saccades suggests that, in both cases, a drift diffusion process can be used to  
 855 keep time and thereby specify the duration of different behavioral modes. The parameters of  
 856 that process can then be used to adjust these durations based on other kinematic properties,  
 857 or based on experience.

### 858 3.2.2 Learning the parameters of a generative model based on experience

859 Given a parameterized generative model of behavior, or behavioral policy  $\pi(a; \vec{\psi})$ , we can now  
 860 ask how the parameters  $\vec{\psi}$  of that policy *should* be modified over time based on experience.  
 861 To this end, we can construct an agent that comes embodied with this generative model, and  
 862 the study how the behavior of that agent changes as we update the generative model based  
 863 on experience (Figure 22).

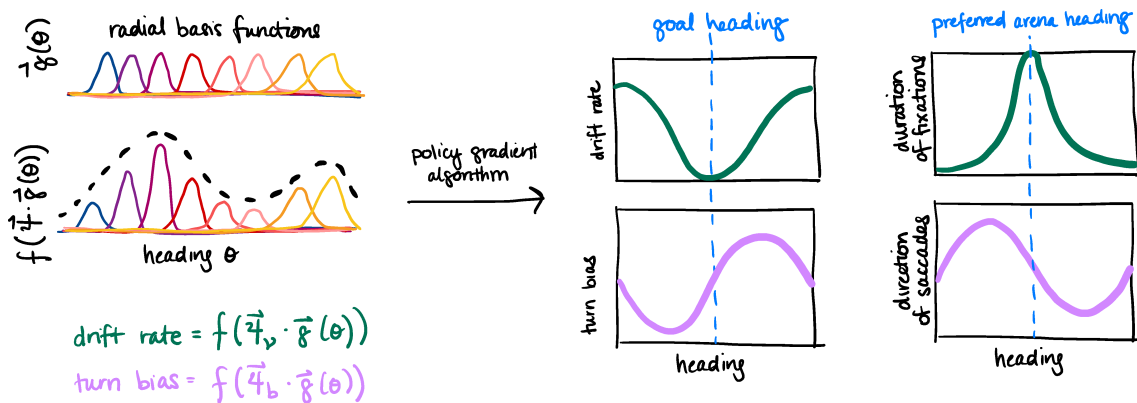


Figure 22: We can use reinforcement learning methods to learn an optimal behavioral policy, and compare it to the behavior of real flies (adapted from [93]).

864 In problems of spatial navigation, a common approach is to discretize the environment into  
 865 states  $\theta$ , and construct a lookup table  $q(\vec{\psi}, \theta)$  that specifies the optimal parameter settings  
 866 for each state. In our setting, these states correspond to different orientations of the agent.  
 867 Moreover, since we are modeling an agent that uses a neural compass to maintain an internal  
 868 representation of its own orientation, we will take these states to correspond to the compass  
 869 heading  $\theta_c$  (see Section 3.1 for a discussion of the compass heading). The agent then uses  
 870 the current parameter settings, together with its own compass heading, to select actions, and  
 871 then iteratively updates the parameters based on the outcomes of those actions in the current  
 872 environment. This process can be quite slow, because there is typically no *a priori* structure  
 873 built into the lookup table, and thus the agent has to experience the consequences of each  
 874 parameter setting at each compass heading in order to make adjustments to these parameters  
 875 (see Appendix D for more details).

876 To speed up this process, the agent can instead learn a continuous function that specifies  
 877 the parameter settings over space. For example, rather than specifying parameter settings  
 878 at each of a set of discretized compass headings, one can use a discrete set of continuous  
 879 basis functions to tile the space of headings, and learn a discrete set of weights on those basis  
 880 functions (Figure 22). This can speed up the learning process because any parameter changes  
 881 that are learned at one heading are immediately used to update parameters at nearby headings  
 882 (if the basis functions are localized in space).

883 We can use this to learn a policy with the following structure:

$$v(\theta_c; \vec{\psi}_v) = f_v(\vec{\psi}_v \cdot \vec{g}(\theta_c)) \quad (35)$$

$$b(\theta_c; \vec{\psi}_b) = f_b(\vec{\psi}_b \cdot \vec{g}(\theta_c)) \quad (36)$$

884 where  $\nu$  and  $\mathbf{b}$  specify the drift rate of fixations and turn bias of saccades derived in the previous  
 885 section, and  $\vec{\mathbf{g}}(\theta_c)$  specifies a set of basis functions that tile compass headings  $\theta_c$ . Informed  
 886 by what we know about the fly’s internal compass, we choose a set of 8 von Mises functions;  
 887 the shape of this function closely matches the shape of the compass bump, and the number  
 888 mimics the known discretization of the compass (see Section 3.1). We can then use a common  
 889 policy-gradient algorithm to update the policy parameters based on experience (see Appendix  
 890 E for a derivation of this algorithm):

$$\begin{aligned}
 &\text{sample action from policy} \\
 &\mathbf{a} \sim \pi(\cdot|\theta_c; \vec{\psi}) \\
 &\text{take action, gather reward} \\
 &\mathbf{R} = r(\mathbf{a}, \theta_c) \\
 &\text{update parameters} \\
 &\Delta \vec{\psi} = \alpha \mathbf{R} \nabla_{\vec{\psi}} \log(\pi(\mathbf{a}|\theta_c; \vec{\psi})) \\
 &\text{update state } \theta_c
 \end{aligned} \tag{37}$$

891 where  $r(\mathbf{a}, \theta_c)$  is a reward function that we will discuss shortly. Using this algorithm, this  
 892 agent can be trained through experience to reliably gather rewards or avoid punishments by  
 893 iteratively updating the parameters of its behavioral policy. We can then compare the output  
 894 of this trained agent to the behavior of real flies after they have gone through the learning  
 895 paradigm schematized in Figure 20.

896 In this paradigm, different orientations of the visual scene are paired with rewards or  
 897 punishments. The visual scene itself consists of a set of 4 visual patterns, each separated by  
 898  $90^\circ$ . The scene was chosen to have a two-fold symmetry, such there are two distinct visual  
 899 patterns,  $\mathbf{A}$  and  $\mathbf{B}$ , each repeated twice (i.e.,  $\mathbf{ABAB}$ ). One set of patterns (either  $\mathbf{A}$  or  $\mathbf{B}$ ) is  
 900 paired with an aversive heat punishment that is delivered by pinging the fly in the abdomen  
 901 with an infrared laser. Thus, if pattern  $\mathbf{A}$  is punished, the fly would experience heat whenever  
 902 it orients within  $\pm 45^\circ$  of the center of pattern  $\mathbf{A}$  (we refer to this region as the “danger zone”,  
 903 and the unpunished region as the “safe zone”). Because we do not know whether learning is  
 904 driven by the negative effects of punishment or the positive effects of relief from punishment,  
 905 we treat the former as generating a negative reward ( $\mathbf{R} = -1$ ), and the latter as generating  
 906 a positive reward ( $\mathbf{R} = +1$ ; note that one can use techniques from inverse reinforcement  
 907 learning to try to deduce this reward function from observed behavior [99]).

908 When we train an artificial agent in this same paradigm, the agent learns to generate  
 909 fixations and saccades who drift rates  $\nu(\theta_c; \vec{\psi}_\nu)$  and turn biases  $\mathbf{b}(\theta_c; \vec{\psi}_b)$  vary approximately  
 910 sinusoidally with the compass heading  $\theta_c$  (Figure 22):

$$\begin{aligned}
 \nu(\theta_c; \vec{\psi}_\nu) &\propto (1 - \cos(\theta_c - \theta_{\text{safe}}))/2 \\
 \mathbf{b}(\theta_c; \vec{\psi}_b) &\propto (1 - \sin(\theta_c - \theta_{\text{safe}}))/2
 \end{aligned} \tag{38}$$

911 where  $\theta_{\text{safe}}$  denotes the center of the safe zone. Low drift rates lead to long fixations at the  
 912 center of the safe zone; a high turn bias to the left of the safe zone generates directed turns  
 913 toward safety.

914 When we compare this learned behavior to the behavior of real flies, we see that flies  
 915 exhibit similar structure to their behavior (to see this, we must align the fixations and saccades  
 916 of individual flies to heading at which these actions were initiated relative to the center of the  
 917 safe zone, and then average across flies; signatures of this structure can also be seen on  
 918 an individual fly basis). However, this behavioral structure is apparent in flies even before  
 919 they experience any laser punishment, provided that we align the behavior to the preferred  
 920 heading of each fly prior to averaging. The fact that both naive and trained flies exhibit the

921 same behavioral structure suggests that this is not something that needs to be learned in a new  
 922 environment; rather, this behavioral structure seems to innately guide how flies sample their  
 923 surroundings. If this is the case, learning need only shift and scale this behavioral structure,  
 924 rather than build it up from scratch.

925 Such a formulation—in which the behavioral policy has a fixed functional form, but flexible  
 926 parameters—dramatically changes the learning process. Rather than building an agent that  
 927 learns individual associations between policy parameters and locations, this suggests that the  
 928 agent should come pre-equipped with a policy that specifies functional relationships between  
 929 policy parameters at all locations in the environment. By building this sort of function into  
 930 the agent, any changes that are learned at one location can be used to update behavior at  
 931 all locations (rather than being restricted to the immediate vicinity of the agent). This can  
 932 significantly speed up the learning process, but comes at the cost of restricting the space of  
 933 behavioral policies that the agent can learn. Moreover, for this to work in real brains, neural  
 934 circuits must be able to enforce this behavioral structure while flexibly shifting and scaling this  
 935 structure to direct movements to different parts of the environment. In the next section, we  
 936 discuss how the brain might achieve this.

### 937 3.2.3 Building a structured behavioral policy

938 To construct the policy given in Eq. 38, we use simple operations that could feasibly be per-  
 939 formed by neural circuits (Figure 23). We assume that the fly maintains encodes its current  
 940 compass heading in a sinusoidal activity profile  $f_c(\theta; \theta_c) = \cos(\theta - \theta_c)$ , where  $\theta_c$  denotes the  
 941 current compass heading, and  $\theta$  describes an anatomical axis that encodes orientation. Note  
 942 that this profile takes on a von Mises shape in the brain region that first encodes the com-  
 943 pass heading, but is reformatted into a sinusoidal shape before propagating to downstream  
 944 brain regions [62]. This sinusoidal shape is thought to support a range of different vector  
 945 computations [100].

946 We assume that the fly maintains a goal activity profile  $f_g(\theta)$  that can take on an arbitrary  
 947 shape. When we convolve these two activity profiles— $f_c(\theta; \theta_c)$  and  $f_g(\theta)$ —the output is  
 948 sinusoidal and peaks when the compass heading is aligned with the circular mean of the goal  
 949 activity profile (Figure 23; see Appendix F for brief derivation). We can thus define the goal  
 950 heading,  $\theta_g$ , to be the circular mean of the goal activity profile  $f_g(\theta)$ .

951 If we use the output of this convolution as a motor drive, this operation will guarantee  
 952 that the motor drive is strongest when the compass heading is aligned with the goal heading.  
 953 However, the profiles given in Eq. 38 require motor drives that are peaked  $90^\circ$  and  $180^\circ$   
 954 from the goal heading. This can be achieved by phase-shifting the compass heading (and thus  
 955 the compass activity profile) before convolving it with the goal activity profile. For a phase  
 956 shift of  $\phi$ , the resulting motor drive will be peaked when the compass heading is aligned with  
 957  $\theta_g - \phi$  (Appendix F). To create the profiles in Eq. 38, we need phase shifts:

$$\begin{aligned}\phi_\nu &= 180^\circ, \\ \phi_b &= 90^\circ.\end{aligned}\tag{39}$$

958 This formulation guarantees that the motor drive will be sinusoidally structured, regardless of  
 959 the shape of the goal activity profile, and leads to a policy of the form:

$$\begin{aligned}\nu(\theta_c; \vec{\psi}_\nu) &\propto A_g \cos(\theta_c - \theta_g + \phi_\nu) + B_\nu, \\ b(\theta_c; \vec{\psi}_b) &\propto A_g \cos(\theta_c - \theta_g + \phi_b) + B_b\end{aligned}\tag{40}$$

960 where  $A_g$  and  $\theta_g$  are the strength and orientation of the circular mean of  $f_g(\theta)$ , respectively,  
 961 and  $(B_\nu, B_b)$  are offsets. This formulation means that the goal activity profile can be updated




Figure 23: Convolution of a sinusoidal compass heading with an arbitrary goal profile leads to a structured motor output (adapted from [93]; see also [101]).

962 over time based on experience, without disrupting the motor drive profile. We can thus use  
 963 simple Hebbian plasticity rules (analogous to those presented in Section 3.1.3) to update the  
 964 goal profile over time:

$$\Delta f_g(\theta, \theta_c; f_g) = \alpha_g \Delta_g \quad (41)$$

965 where

$$\Delta_g = \begin{cases} +[f_c(\theta; \theta_c) - f_g(\theta)]_+ \Theta(1 - f_g) - [f_g(\theta) - f_c(\theta; \theta_c)]_+ \Theta(f_g) & R(\theta) > 0 \\ -[f_c(\theta; \theta_c) - f_g(\theta)]_+ \Theta(f_g) + [f_g(\theta) - f_c(\theta; \theta_c)]_+ \Theta(1 - f_g) & R(\theta) < 0 \\ 0 & R(\theta_A) = 0 \end{cases} \quad (42)$$

966 Here,  $[\cdot]_+$  denotes rectification, and  $\Theta(\cdot)$  is the heaviside function. This plasticity rule will  
 967 strengthen the goal profile at the current compass heading when the fly is being rewarded  
 968 ( $R > 0$ ), and will weaken the goal profile at the current compass heading when the fly is being  
 969 punished ( $R < 0$ ).

970 The fly central brain has circuit motifs that could plausibly implement this set of computa-  
 971 tions [63, 93] (Figure 24). As discussed in Section 3.1, the compass heading is maintained by a  
 972 ring attractor network in a region of the brain called the ellipsoid body (EB). From the EB, the  
 973 compass activity profile is reformatted from a von-Mises-like shape to a sinusoidal shape [62]  
 974 before traveling to a downstream region called the fan-shaped body (FB) that is thought to  
 975 carry and combine internal representations about the fly's goals and future actions [63]. There  
 976 are two populations of neurons—so-called PFL2 and PFL3 populations—that are thought to  
 977 implement the phase shifts in Eq. 39, respectively, and that project either bilaterally or uni-  
 978 laterally to pre-motor regions that are involved in initiating rightward and leftward turns.  
 979 The PFL2 population implements a  $180^\circ$  phase shift and projects bilaterally to both premotor  
 980 regions, and is thus well-positioned to control straight flight. The two PFL3 populations im-  
 981 plement  $\pm 90^\circ$  phase shifts and project unilaterally to one or the other premotor regions, and  
 982 are thus well-positioned to control directed turns. Finally, there are dopaminergic neurons  
 983 that synapse onto both PFL2 and PFL3 populations, and that themselves receive inputs that  
 984 are shaped by the fly's current compass heading; thus, these neurons are well-positioned to



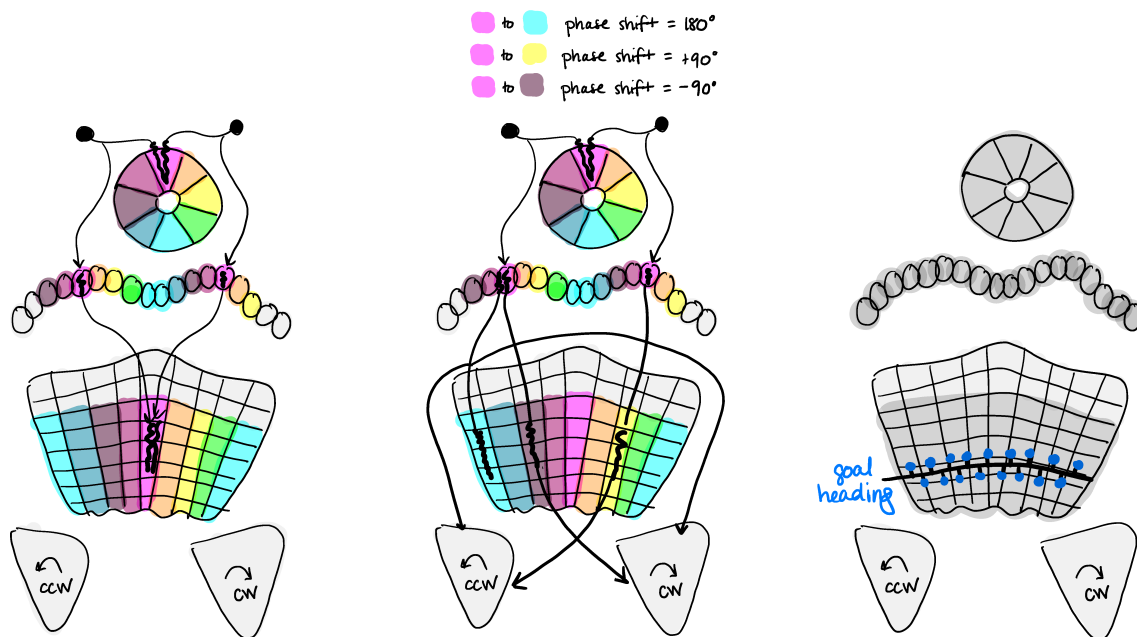


Figure 24: Left: Neurons in the fly central brain implement an anatomical reference frame tethered to the fly's current heading. Center: additional populations of neurons implement anatomical phase shifts in this reference frame. Right: putative goal neurons could store and update goal headings (adapted from [63] and [93]).

985 update the strength of goal synapses in a heading-dependent manner, which could in turn be  
 986 used to drive behavior through the PFL populations.

### 987 3.3 Coupling sensory representations and behavior in closed loop

988 In Section 3.1, we saw how small neural circuits could maintain and accurately update an  
 989 internal representation of heading, and could tether this representation to sensory cues in  
 990 the environment through a form of unsupervised learning. In Section 3.2, we saw how this  
 991 internal representation of heading could be used to update an internal representation of goals  
 992 in the environment through a form of reinforcement learning. These two systems—learning  
 993 where you are in your environment, and learning goals in that environment—are coupled  
 994 through the same internal representation of heading, and through the actions that are selected  
 995 by the behavioral policy that is tethered to this representation (Figure 25). As a result, any  
 996 inaccuracies in this representation will impact the formation of goals, the selection of actions,  
 997 and the subsequent sensory stimuli that are used to update this representation. As a result,  
 998 the coupling of these systems can shape individual variability in behavior and learning.

999 This becomes apparent if we revisit the ideas from Section 3.1.3, about how the internal  
 1000 representation of heading becomes tethered to the outside world. At the end of that section, we  
 1001 introduced the idea that symmetries in the visual scene can lead to a scenario in which plasticity  
 1002 tethers the bump to one of two symmetric location in the environment. If two ring neurons are  
 1003 active for a particular view of a symmetric scene, but one strongly inhibits the bump and the  
 1004 other only weakly inhibits the bump, this induces competition between these two ring neurons,  
 1005 and the bump will tend to jump to the location of weakest inhibition. In other words, the  
 1006 compass will develop a two-to-one mapping, in which two different symmetric orientations of  
 1007 the scene will be mapped onto a single compass heading. This type of confusion is analogous  
 1008 to the confusion you might expect to experience inside a room with two identical doors on  
 1009 opposite sides of the room. When you initially enter the room, you might remember which

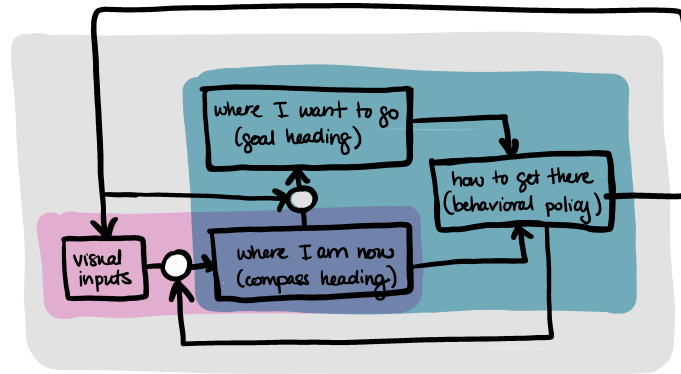


Figure 25: Two learning systems interact to guide behavior. One learning system (pink) maps visual inputs onto an internal compass heading; a second learning system (green) uses the compass heading to learn a goal heading and guide actions via a behavioral policy. Plasticity acts at specific locations (white circles) to allow these representations to change over time based on experience.

1010 of the two doors you used to enter, but after some time, you might get confused about which  
 1011 door is which. We can see this confusion in the dynamics of the compass bump, and we can  
 1012 watch the compass bump jump between different orientations that correspond to symmetric  
 1013 views of the visual world [93].

1014 This highlights the important difference between the absolute properties of the visual en-  
 1015 vironment and the relative nature of the internal constructs that we build to represent that  
 1016 environment. In many cases, these are closely matched, but in some cases, our internal per-  
 1017 ceptions differ markedly from the outside world. This has interesting implications for any  
 1018 downstream behavior that is tethered to these representations. For example, if the fly's ac-  
 1019 tions are tethered to this internal representation, as evidence suggests, then its actions will  
 1020 follow this representation, even as that representation is jumping over time. This can be ad-  
 1021 vantageous in an environment where rewards and punishments are coupled to these same  
 1022 symmetries of the environment. For example, given a policy that is tethered to the difference  
 1023 between the fly's current and goal headings (Eq. 40), a jumping heading bump will serve to  
 1024 copy over this policy at the arena headings that correspond to symmetric views of the visual  
 1025 scene. The means that the fly need only learn one goal heading; its neural circuits will effi-  
 1026 ciently structure behavior about that goal heading, and its jumping compass bump will copy  
 1027 that structure at multiple locations that share the same visual patterns and reward structures.

1028 For this to work effectively, the mapping of the visual world onto the compass must be  
 1029 stabilized, which takes time in a new environment. As a result, stronger initial goal headings,  
 1030 which lead to more exploitative behavior about the goal heading, can help to quickly anchor  
 1031 the mapping of the visual world onto the compass, particular in the presence of visual sym-  
 1032 metries. This, in turn, can help the fly more quickly update its goal heading in the face of  
 1033 rewards and punishments. Weak goal headings, on the other hand, can lead to more diffusive,  
 1034 exploratory behavior, can slow the development of the visual map in the presence of symme-  
 1035 tries, and can lead to misalignment between the learning of the visual map and the learning  
 1036 of goals tethered to that visual map. This, in turn, can slow the overall learning process. As a  
 1037 result, individual variability in learning can arise purely because of the dynamics that couple  
 1038 these multiple internal representations. See [93] for further discussion.

### 1039 3.4 Summary

1040 In this section, we explored some of the challenges of building compact internal represen-  
1041 tations of the outside world, bootstrapping these representations by using one evolving rep-  
1042 resentation to update another, and ultimately tethering efficient behavior to these internal  
1043 representations. One of the major insights from this line of work is that pre-motor circuits in  
1044 the brain seems to employ hardwired anatomical motifs for exploiting structure in an animal's  
1045 predictable relationship to its surroundings, just as early sensory circuits exploit predictable  
1046 structure in incoming sensory signals. In the field of machine learning, these architectural  
1047 constraints are often called 'structural priors' or 'inductive biases', and ongoing work in the  
1048 field seeks to understand what sorts of inductive biases can lead to more flexible and gener-  
1049 alizable architectures [102, 103]. Here, we saw how inductive biases in the fly brain—in the  
1050 form of anatomical motifs that efficiently structure behavior about a single goal heading—can  
1051 significantly speed learning by reducing the number of associations that an animal has to ex-  
1052 perience. This suggests an exciting direction for future research, in terms of understanding  
1053 how the brain should best balance the speed enabled by hardwired circuits, and the flexibility  
1054 afforded by plasticity within these circuits. More broadly, this prompts the question as to what  
1055 additional features of an animal's relationship to its environment are predictable enough to  
1056 warrant hard-wired solutions—something that will be aided by quantitative analyses of natural  
1057 behavior in the context of ethologically-relevant tasks.

## 1058 4 Outlook

1059 In these notes, we have touched upon different theoretical frameworks—including efficient  
1060 coding, Bayesian inference, and reinforcement learning—that are used to understand sensory  
1061 coding, inference, and action selection. We then used aspects of these frameworks in the  
1062 context of a specific example—navigational learning in the fruit fly—to understand how real  
1063 brains make sense of incoming sensory stimuli, build internal representations of their relation-  
1064 ship to the external world, and use these internal representations to guide behavior. In both of  
1065 these contexts, we discussed how these frameworks provide a perspective on not only *how* the  
1066 brain operates, but also *why* it might operate in one manner versus another. Looking ahead,  
1067 one major opportunity in theoretical neuroscience is to understand how these different frame-  
1068 works relate to one another, and how they inform and constrain one another. We saw how  
1069 resource constraints, as considered in the framework of efficient coding, can impact the speed  
1070 and accuracy of Bayesian inference, and how constraints on the precision of inference can  
1071 in turn impact action selection. We also saw how behavioral policies can exploit the reliable  
1072 statistics of movement to speed learning, just as sensory codes can exploit the reliable statistics  
1073 of sensory stimuli to speed inference. Finally, we saw that it is becoming possible to test the  
1074 predictions of these interconnected frameworks in real brains, where the fruit fly—with its  
1075 targeted genetic tools, near-complete connectome, and behavioral repertoire—provides one  
1076 exciting test bed to link multiple successive computations in closed-loop.

## 1077 Acknowledgements

1078 Wiktor Młynarski and Marcella Noorman provided helpful feedback on these notes. As in-  
1079 dicated in the respective section headings, some of the work and perspectives presented in  
1080 Section 2 derived from collaborations with Wiktor Młynarski and Tzuhsuan (Maz) Ma; sim-  
1081 ilarly, some of the work and perspectives presented in Section 3 derived from collaborations

1082 with Marcella Noorman, Brad Hulse, Vivek Jayaraman, Sung Soo Kim, Larry Abbott, Sandro  
1083 Romani, and Chuntao Dan. Some perspectives presented in the introduction derived from  
1084 collaborations with Gašper Tkačik, Jonathan Victor, Mary Conte, John Briguglio, and Vijay  
1085 Balasubramanian.

1086 **Funding information** This work was supported by the Howard Hughes Medical Institute.

## 1087 A Decomposing mutual information

1088 The mutual information between  $X$  and  $Y$  is defined as:

$$I(X; Y) = \sum_{x \in X, y \in Y} P(x, y) \log \left( \frac{P(x, y)}{P(x)P(y)} \right) \quad (\text{A.1})$$

1089 We can write  $P(x, y) = P(x|y)P(y)$ , which allows us to re-express mutual information:

$$\begin{aligned} I(X; Y) &= \sum_{x \in X, y \in Y} P(x|y)P(y) \log \left( \frac{P(x|y)P(y)}{P(x)P(y)} \right) \\ &= \sum_{y \in Y} P(y) \sum_{x \in X} P(x|y) [\log P(x|y) - \log P(x)] \end{aligned} \quad (\text{A.2})$$

1090 Noting that the conditional entropy of  $X$  given  $Y$  is defined as:

$$H(X|Y) = \sum_{y \in Y} P(y) \left[ - \sum_{x \in X} P(x|y) \log P(x|y) \right], \quad (\text{A.3})$$

1091 we can use this to rewrite the mutual information as:

$$\begin{aligned} I(X; Y) &= -H(X|Y) - \sum_{x \in X, y \in Y} P(x, y) \log P(x) \\ &= -H(X|Y) - \sum_{x \in X, y \in Y} P(y|x)P(x) \log P(x) \\ &= -H(X|Y) - \sum_{x \in X} P(x) \log P(x) \sum_{y \in Y} P(y|x) \\ &= -H(X|Y) + H(X) \end{aligned} \quad (\text{A.4})$$

1092 where we have used the fact that  $H(X) = - \sum_{x \in X} P(x) \log P(x)$ , and  $\sum_{y \in Y} P(y|x) = 1$ .

## 1093 B Maximum entropy distributions

1094 Consider a neuron that can produce a discrete set of discriminable responses. The entropy of  
1095 the response distribution is:

$$H = - \sum_i p_i \log(p_i) \quad (\text{B.1})$$

1096 subject to the constraint  $\sum_i p_i = 1$ .

1097 To find the entropy-maximizing distribution of responses, we can define the Lagrangian  $\mathcal{L}$   
1098 with Lagrange multiplier  $\lambda_0$ :

$$\mathcal{L}(p, \lambda_0) = - \sum_i p_i \log(p_i) + \lambda_0 \left( \sum_i p_i - 1 \right) \quad (\text{B.2})$$

1099 Maximizing this with respect to  $p_i$  gives:

$$\begin{aligned} \frac{\partial \mathcal{L}(p, \lambda_0)}{\partial p} = 0 &\implies -\log(p_i) - 1 + \lambda_0 \\ &\implies p_i = \exp(\lambda_0 - 1) \end{aligned} \quad (\text{B.3})$$

1100 Imposing the constraint  $\sum_i p_i = 1$  means:

$$\begin{aligned}\sum \exp(\lambda_0 - 1) &= 1 \\ \implies N \exp(\lambda_0 - 1) &= 1 \\ \implies p_i &= 1/N\end{aligned}\tag{B.4}$$

1101 Intuitively, given a constraint on the number of discriminable response levels, the maximum  
1102 entropy distribution is flat (and corresponds to ‘histogram equalization’).

1103 If we now consider a continuous distribution of responses but impose a constraint on the  
1104 mean firing rate  $\mu$ , this gives:

$$\mathcal{L}(p, \lambda_0, \lambda_1) = - \int p(r) \log(p(r)) dr + \lambda_0 \left( \int p(r) dr - 1 \right) + \lambda_1 \left( \int r p(r) dr - \mu \right)\tag{B.5}$$

1105 and

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial p} &= -\log(p_i) - 1 + \lambda_0 + \lambda_1 r = 0 \\ \implies p(r) &= \exp(\lambda_0 - 1) \exp(\lambda_1 r).\end{aligned}\tag{B.6}$$

1106 Plugging this back into the equations for our constraints, we have:

$$\begin{aligned}\lambda_1 &= -1/\mu; \\ \exp(\lambda_0 - 1) &= 1/\mu\end{aligned}$$

1107 which gives

$$p(r) = \frac{1}{\mu} \exp\left(-\frac{r}{\mu}\right)\tag{B.7}$$

1108 Thus, given a constraint on the mean firing rate, the distribution of neural responses is expo-  
1109 nential. It is straightforward to use the same logic to show that an additional constraint on  
1110 the variance of neural responses leads to a Gaussian distribution.

## 1111 C Bayesian ideal observer model in a switching environment

1112 Here, we derive the ideal observer model for the scenario described in Section 2.2, which must  
1113 infer the current state  $\theta_t$  of the environment given the history of past sensory signals  $s_{\tau \leq t}$ .

1114 Given random variables  $A$  and  $B$ , Bayes Rule states:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}\tag{C.1}$$

1115 For three variables  $A, B, C$ , this becomes:

$$P(A|B, C) = \frac{P(B|A, C)P(A|C)}{P(B|C)}\tag{C.2}$$

1116 We can now make the substitutions:  $A = \theta_t$ ,  $B = s_t$ , and  $C = s_{\tau < t}$ :

$$P(\theta_t | s_t, s_{\tau < t}) = \frac{P(s_t | \theta_t, s_{\tau < t}) P(\theta_t | s_{\tau < t})}{P(s_t | s_{\tau < t})}\tag{C.3}$$



1117 This distribution must be normalized, which gives:

$$\sum_{\theta_t} P(\theta_t | s_t, s_{\tau < t}) = 1 = \sum_{\theta_t} \frac{P(s_t | \theta_t, s_{\tau < t}) P(\theta_t | s_{\tau < t})}{P(s_t | s_{\tau < t})} \quad (\text{C.4})$$

1118 Since the distribution in the denominator does not depend on  $\theta_t$ , we can pull it out of the sum  
1119 and define this to be the normalization constant  $\Omega$ :

$$\Omega \equiv P(s_t | s_{\tau < t}) = \sum_{\theta_t} P(s_t | \theta_t, s_{\tau < t}) P(\theta_t | s_{\tau < t}) \quad (\text{C.5})$$

1120 We now have:

$$P(\theta_t | s_t, s_{\tau < t}) = \frac{1}{\Omega} P(s_t | \theta_t, s_{\tau < t}) P(\theta_t | s_{\tau < t}) \quad (\text{C.6})$$

1121 The second distribution can be expanded and simplified as follows:

$$P(\theta_t | s_{\tau < t}) = \sum_{\theta_{t-1}} P(\theta_t | \theta_{t-1}, s_{\tau < t}) P(\theta_{t-1} | s_{\tau < t}) \quad (\text{C.7})$$

$$= \sum_{\theta_{t-1}} P(\theta_t | \theta_{t-1}) P(\theta_{t-1} | s_{t-1}, s_{\tau < t-1}) \quad (\text{C.8})$$

1122 Putting this together, we have:

$$P(\theta_t | s_t, s_{\tau < t}) = \frac{1}{\Omega} P(\theta_t | s_t) \sum_{\theta_{t-1}} P(\theta_t | \theta_{t-1}) P(\theta_{t-1} | s_{t-1}, s_{\tau < t-1}) \quad (\text{C.9})$$

1123 For the specific case of a two-state environment, we can define  $P_t^L \equiv P(\theta_t = \theta^L | s_t, s_{\tau < t})$ ,  
1124 and  $P_t^H = (\mathbf{1} - P_t^L)$ . The posterior distribution is thus fully specified by a single number: the  
1125 probability that the environment is in the low state, given the history of past observations. In  
1126 matrix notation, we now have:

$$\begin{bmatrix} P_t^L \\ \mathbf{1} - P_t^L \end{bmatrix} = \frac{1}{\Omega} \begin{bmatrix} P(s_t | \theta_t = \theta^L) & P(s_t | \theta_t = \theta^H) \end{bmatrix} \begin{bmatrix} 1 - p_s & p_s \\ p_s & 1 - p_s \end{bmatrix} \begin{bmatrix} P_{t-1}^L \\ \mathbf{1} - P_{t-1}^L \end{bmatrix} \quad (\text{C.10})$$

1127 Or equivalently,

$$P_t^L = \frac{1}{\Omega} P(s_t | \theta_t = \theta^L) [(1 - p_s) P_{t-1}^L + p_s (1 - P_{t-1}^L)] \quad (\text{C.11})$$

## 1128 D Primer on reinforcement learning

1129 This is meant to provide a quick-and-dirty tutorial on reinforcement learning; for a more thor-  
1130 ough overview and background, please refer to Sutton & Barto [40], which provided the source  
1131 for these notes.

1132 Reinforcement learning considers the process of learning to select appropriate actions  $\mathbf{a}$  in  
1133 particular states  $\theta$  of the environment, based on feedback from the environment in the form  
1134 of rewards  $r$ . Note that it is typical to denote states with the variable  $\mathbf{s}$ , instead of  $\theta$ ; here we  
1135 use  $\theta$  to highlight that these states are typically hidden from the agent must be inferred from  
1136 sensory signals, which we denote  $\mathbf{s}$ . In Section 2.2, we used such latent states to parameterize  
1137 distributions of sensory stimuli; in Section 2.3, the latent state specified the reward probability  
1138 of different levers that could be sampled by an agent. And in Section 3, the latent state specified  
1139 the animal's heading in the environment, which must be inferred from visual and self-motion  
1140 signals.

1141 The goal of a reinforcement learning agent is to maximize long-term accumulated reward,  
 1142 which is called the return  $G_t$  and is defined as the total sum of all rewards received at future  
 1143 times  $t' > t$ :

$$G_t \equiv R_{t+1} + R_{t+2} + R_{t+3} + \dots \quad (\text{D.1})$$

1144 As written, this expression assumes that all future rewards are equally important. To account  
 1145 for the fact the immediate rewards might be more valuable than rewards in the distant future,  
 1146 we can “discount” future rewards with a discount factor  $\gamma \in [0, 1]$ :

$$\begin{aligned} G_t &= R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots \\ &= \sum_{t'=\ell+1}^{\infty} \gamma^{t'-1} R_{t'} \\ &= R_{t+1} + \gamma G_{t+1} \end{aligned} \quad (\text{D.2})$$

1147 where  $\gamma = 1$  defines a far-sighted agent that values all rewards equally, and  $\gamma = 0$  defines a  
 1148 purely myopic agent that only values immediate rewards.

1149 The return  $G_t$  represents the sum of actual future rewards, which is not knowable to any  
 1150 agent. Instead, an agent can compute the *expected* future rewards, starting in state  $\theta$  and  
 1151 following a policy  $\pi(\mathbf{a}|\theta)$ . The policy specifies the probability of selecting a given action  $\mathbf{a}$   
 1152 from a state  $\theta$ . If it is helpful to have a concrete example in mind, consider a scenario in which  
 1153 an agent navigates in 2D environment to gather rewards that are given when the agent reaches  
 1154 specific locations in the environment. A typical way to treat this scenario is to discretize the  
 1155 environment into a grid of locations; these locations then serve as our states (so, in a  $10 \times 10$   
 1156 grid, there are 100 states). At any grid location, it is often assumed that there are at most 5  
 1157 available actions: stay, move 1 step up, move 1 step down, move 1 step left, or move 1 step  
 1158 right. A policy can then be specified as the probability of taking each of the 5 actions from  
 1159 each of the 100 states in the environment.

1160 With this example in mind, we can start by computing the expected future rewards that  
 1161 will be obtained one timestep in the future:

$$\mathbb{E}_\pi[R_{t+1}|\theta_t = \theta] = \sum_{\mathbf{a}} \pi(\mathbf{a}|\theta) \sum_{\theta'} p(\theta'|\theta, \mathbf{a}) r(\theta, \mathbf{a}, \theta') \quad (\text{D.3})$$

1162 where  $p(\theta'|\theta, \mathbf{a})$  governs the dynamics of the environment and specifies the probability that  
 1163 the agent transitions to a state  $\theta'$  when beginning in state  $\theta$  and taking action  $\mathbf{a}$ ; the reward  
 1164 function  $r(\theta', \mathbf{a}, \theta)$  specifies the reward received under the same transition. We can use Eq.  
 1165 D.3 to define a state-value function  $v_\pi(\theta)$  that defines how ‘good’ it is (in terms of expected  
 1166 return) to begin in state  $\theta$  and follow the policy  $\pi$ :

$$v_\pi(\theta) \equiv \mathbb{E}_\pi[G_t | \theta_t = \theta] \quad (\text{D.4})$$

1167 Using the result from Eqs. D.2-D.3, we can write this as:

$$\begin{aligned} v_\pi(\theta) &= \mathbb{E}_\pi[R_{t+1} + \gamma G_{t+1} | \theta_t = \theta] \\ &= \sum_{\mathbf{a}} \pi(\mathbf{a}|\theta) \sum_{\theta'} p(\theta'|\theta, \mathbf{a}) \left[ r(\theta, \mathbf{a}, \theta') + \gamma \mathbb{E}_\pi[G_{t+1} | \theta_{t+1} = \theta'] \right] \\ &= \sum_{\mathbf{a}} \pi(\mathbf{a}|\theta) \sum_{\theta'} p(\theta'|\theta, \mathbf{a}) \left[ r(\theta, \mathbf{a}, \theta') + \gamma v_\pi(\theta') \right] \end{aligned} \quad (\text{D.5})$$

1168 This is the recursive Bellman equation for the state-value function [104]. We can also use this  
 1169 to define a state-action value function  $q_\pi(\theta, \mathbf{a})$  that defines the value of starting in a state  $\theta$ ,  
 1170 taking an action  $\mathbf{a}$ , and following the policy  $\pi$  from then onward:

$$q_\pi(\theta, \mathbf{a}) \equiv \sum_{\theta'} p(\theta'|\theta, \mathbf{a}) \left[ r(\theta, \mathbf{a}, \theta') + \gamma v_\pi(\theta') \right] \quad (\text{D.6})$$

1171 for which there is an analogous Bellman equation:

$$q_\pi(\theta, \mathbf{a}) = \sum_{\theta'} p(\theta'|\theta, \mathbf{a}) \left[ r(\theta, \mathbf{a}, \theta') + \gamma \sum_{\mathbf{a}'} \pi(\mathbf{a}'|\theta') q_\pi(\theta', \mathbf{a}') \right] \quad (\text{D.7})$$

1172 The optimal value function satisfies the Bellman optimality equation:

$$\begin{aligned} v_*(\theta) &= \max_{\mathbf{a}} \mathbb{E}_\pi [R_{t+1} + \gamma v^*(\theta_{t+1}) | \theta_t = \theta, \mathbf{a}_t = \mathbf{a}] \\ &= \max_{\mathbf{a}} \sum_{\theta'} p(\theta'|\theta, \mathbf{a}) [r(\theta, \mathbf{a}, \theta') + \gamma v_*(\theta')] \end{aligned} \quad (\text{D.8})$$

1173 and analogously for  $q_*(\theta, \mathbf{a})$ .

1174 Note that one can use a value function  $v_\pi(\theta)$  or action-value function  $q_\pi(\theta, \mathbf{a})$  to define  
 1175 a policy. For example, a purely greedy policy would select actions that maximize the current  
 1176 value:

$$\mathbf{A} = \operatorname{argmax}_{\mathbf{a}} q(\theta, \mathbf{a}) \quad (\text{D.9})$$

1177 If  $q = q_*$  is the optimal action-value function, then Eq. D.9 defines the optimal policy  
 1178  $\pi_*(\mathbf{a}|\theta) = \operatorname{argmax}_{\mathbf{a}} (q_*(\theta, \mathbf{a}))$ . However, if  $q$  is not optimal, it can be advantageous to use  
 1179 a policy that balances exploitative actions (which maximize the current estimate of  $q$ ) with  
 1180 exploratory actions (which have lower expected value but could result in higher long-term  
 1181 payoffs). A simple version of such a policy is called an ‘epsilon-greedy’ policy, in which the  
 1182 agent chooses the greedy (exploitative) action with probability  $(1 - \epsilon)$ , and chooses a random  
 1183 (exploratory) action with probability  $\epsilon$ , where  $\epsilon \in [0, 1]$  is typically chosen to be small. Al-  
 1184 ternatively, one can tie the degree of exploration to the value function by choosing a ‘softmax’  
 1185 policy:

$$\pi(\mathbf{a}|\theta) \propto \exp(\beta q(\theta, \mathbf{a})) \quad (\text{D.10})$$

1186 where  $\beta \rightarrow \infty$  drives purely exploitative actions that maximize  $q$ , and  $\beta \rightarrow 0$  drives purely  
 1187 random actions.

1188 A goal of reinforcement learning is to use value functions to determine the optimal poli-  
 1189 cies that maximize the return  $G$ . Historically, there have been three main sets of methods  
 1190 for learning value functions. The first are dynamic programming methods, which assume a  
 1191 model of the environment in the form of  $p(\theta'|\theta, \mathbf{a})$  and  $r(\theta, \mathbf{a}, \theta')$ , and learn  $v_*$ ,  $q_*$ , and  
 1192  $\pi_*$  via bootstrapping (i.e., updating an estimate of one quantity based on an estimate of an-  
 1193 other quantity). Value iteration, which we introduced in Section 2.3, is one such method that  
 1194 converts Eq. D.8 into an update rule for the value function:

$$v_{t+1}(\theta) = \max_{\mathbf{a}} \sum_{\theta', r} p(\theta'|\theta, \mathbf{a}) \left[ r(\theta, \mathbf{a}, \theta') + \gamma v_t(\theta') \right] \quad (\text{D.11})$$

1195 The second set of methods are Monte Carlo methods, which do not assume a model of the  
 1196 environment, and instead directly estimate  $v^*$ ,  $q^*$ , and  $\pi^*$  via sampling (i.e., from the outcome

1197 of simulated experiences). An example of a Monte Carlo update is the following:

$$v_{t+1}(\theta) = v_t(\theta) + \alpha[G_t - v_t(\theta)] \quad (\text{D.12})$$

1198 where the return  $G_t$  must be computed from the outcome of simulated experiences.  $G_t$  is then  
1199 used as a target to update  $v_t$ .

1200 A third class of methods, temporal difference (TD) methods, combine aspects of dynamic  
1201 programming and Monte Carlo methods to learn from experience via bootstrapping [40]. The  
1202 simplest TD method replaces the return  $G_t$  in Eq. D.12 with the estimate  $G_t \approx R_{t+1} + \gamma v_t(\theta_{t+1})$ :

$$v_{t+1}(\theta_t) = v_t(\theta_t) + \alpha[R_{t+1} + \gamma v_t(\theta_{t+1}) - v_t(\theta_t)] \quad (\text{D.13})$$

$$q_{t+1}(a_t, \theta_t) = q_t(a_t, \theta_t) + \alpha[R_{t+1} + \gamma q_t(a_{t+1}, \theta_{t+1}) - q_t(a_t, \theta_t)] \quad (\text{D.14})$$

1203 Here,  $\delta_t = R_{t+1} + \gamma v_t(\theta_{t+1}) - v_t(\theta_t)$  (and similarly  $\delta_t = R_{t+1} + \gamma q_t(a_{t+1}, \theta_{t+1}) - q_t(a_t, \theta_t)$ )  
1204 is often referred as the TD error. Eq. D.14 forms the basis of the so-called SARSA learning  
1205 algorithm. Note that the updating of the value functions depends on the actions taken and  
1206 rewards received, which in turn depend on the policy used to select those actions. Thus, when  
1207 using one of the policies defined in Eqs. D.9-D.10, the actions selected via the policy will  
1208 change over time as the value functions are updated via Eqs. D.13-D.14.

1209 With an algorithm like SARSA, the agent can only update the value of states as these states  
1210 are visited. To update states that were visited in the past, we can augment this update rule  
1211 with an eligibility trace  $Z_t(\theta)$  that allows previously-visited to be eligible for updating:

$$v_{t+1}(\theta_t) = v_t(\theta_t) + \alpha[R_{t+1} + \gamma v_t(\theta_{t+1}) - v_t(\theta_t)]Z_t(\theta_t) \quad (\text{D.15})$$

1212 where  $Z_t(\theta)$  follows its own update rule:

$$Z_t(\theta) = \begin{cases} \lambda \gamma Z_{t-1}(\theta) & \theta \neq \theta_t \\ 1 + \lambda \gamma Z_{t-1}(\theta) & \theta = \theta_t \end{cases} \quad (\text{D.16})$$

1213 Here,  $\lambda \in [0, 1]$  is a trace-decay parameter that specifies how quickly the eligibility of each  
1214 state will decay over time. When  $\lambda = 0$ , only the current state can be updated; when  $\lambda = 1$ ,  
1215 the eligibility of an unvisited state falls by  $\gamma$  with each timestep. Larger values of  $\lambda$  cause the  
1216 agent to associate previously-visited states with current rewards.

## 1217 E Derivation of policy gradient algorithm

1218 This follows the derivation given in Sutton & Barto [40]. We begin by defining a performance  
1219 measure  $J(\vec{\psi})$  that we want to maximize, given a policy  $\pi(a|\theta; \vec{\psi})$  parameterized by  $\vec{\psi}$ :

$$J(\vec{\psi}) \equiv v_{\pi_{\vec{\psi}}}(\theta) \quad (\text{E.1})$$

1220 where  $v_{\pi_{\vec{\psi}}}(\theta)$  is the value for a policy  $\pi$ , parameterized by  $\vec{\psi}$ , starting in state  $\theta$ .

1221 We want to find the parameter update that increases performance over time:

$$\begin{aligned} \nabla_{\vec{\psi}} v_{\pi}(\theta) &= \nabla_{\vec{\psi}} \left[ \sum_a \pi(a|\theta) q_{\pi}(\theta, a) \right] \\ &= \sum_a \left[ q_{\pi}(\theta, a) \nabla_{\vec{\psi}} \pi(a|\theta) + \pi(a|\theta) \nabla_{\vec{\psi}} q_{\pi}(\theta, a) \right] \end{aligned} \quad (\text{E.2})$$

1222 The second term in this sum,  $*$ , can be written as:

$$\begin{aligned} * &= \pi(\mathbf{a}|\boldsymbol{\theta}) \nabla_{\vec{\psi}} \sum_{\theta', r} p(\theta', r|\boldsymbol{\theta}, \mathbf{a})(r + v_{\pi}(\theta')) \\ &= \pi(\mathbf{a}|\boldsymbol{\theta}) \nabla_{\vec{\psi}} \sum_{\theta'} p(\theta'|\boldsymbol{\theta}, \mathbf{a}) \nabla_{\vec{\psi}} v_{\pi}(\theta') \end{aligned} \quad (\text{E.3})$$

1223 We can now note that

$$\nabla_{\vec{\psi}} v_{\pi}(\boldsymbol{\theta}) = \sum_{\mathbf{a}} \left[ q_{\pi}(\boldsymbol{\theta}, \mathbf{a}) \nabla_{\vec{\psi}} \pi(\mathbf{a}|\boldsymbol{\theta}) + \pi(\mathbf{a}|\boldsymbol{\theta}) \sum_{\theta'} p(\theta'|\boldsymbol{\theta}, \mathbf{a}) \nabla_{\vec{\psi}} v_{\pi}(\theta') \right] \quad (\text{E.4})$$

1224 and in a similar manner,  $\nabla_{\vec{\psi}} v_{\pi}(\theta')$  can be ‘rolled out’ and expressed in terms of  $\nabla_{\vec{\psi}} v_{\pi}(\theta'')$ :

$$\begin{aligned} \nabla_{\vec{\psi}} v_{\pi}(\boldsymbol{\theta}) &= \sum_{\mathbf{a}} \left[ q_{\pi}(\boldsymbol{\theta}, \mathbf{a}) \nabla_{\vec{\psi}} \pi(\mathbf{a}|\boldsymbol{\theta}) + \pi(\mathbf{a}|\boldsymbol{\theta}) \sum_{\theta'} p(\theta'|\boldsymbol{\theta}, \mathbf{a}) \right. \\ &\quad \left. \sum_{\mathbf{a}'} \left[ q_{\pi}(\theta', \mathbf{a}') \nabla_{\vec{\psi}} \pi(\mathbf{a}'|\theta') + \pi(\mathbf{a}'|\theta') \sum_{\theta''} p(\theta''|\theta', \mathbf{a}') \right. \right. \\ &\quad \left. \left. \nabla_{\vec{\psi}} v_{\pi}(\theta'') \right] \right]. \end{aligned} \quad (\text{E.5})$$

1225 Note that the first two rows of the above equation are identical in form, but are weighted  
1226 by the probability of transitioning between states under policy  $\pi$ . Replacing  $\boldsymbol{\theta}$  with its value  
1227 sampled at time  $t$ , we can write:

$$\begin{aligned} \nabla_{\vec{\psi}} J(\vec{\psi}) &= \mathbb{E}_{\pi} \left[ \sum_{\mathbf{a}} q_{\pi}(\boldsymbol{\theta}_t, \mathbf{a}) \nabla_{\vec{\psi}} \pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi}) \right] \\ &= \mathbb{E}_{\pi} \left[ \sum_{\mathbf{a}} \pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi}) q_{\pi}(\boldsymbol{\theta}_t, \mathbf{a}) \frac{\nabla_{\vec{\psi}} \pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi})}{\pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi})} \right] \\ &= \mathbb{E}_{\pi} \left[ \sum_{\mathbf{a}} \pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi}) q_{\pi}(\boldsymbol{\theta}_t, \mathbf{a}) \nabla_{\vec{\psi}} \log \pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi}) \right] \end{aligned} \quad (\text{E.6})$$

1228 where we have multiplied and divided by  $\pi(\mathbf{a}|\boldsymbol{\theta}_t; \vec{\psi})$  in the second line, and written  $\nabla \mathbf{x} / \mathbf{x} = \nabla \log \mathbf{x}$   
1229 in the third line. If we now replace  $\mathbf{a}$  by a sampled action  $\mathbf{a}_t \sim \pi(\cdot|\boldsymbol{\theta}_t)$ , we have:

$$\nabla_{\vec{\psi}} J(\vec{\psi}) = \mathbb{E}_{\pi} \left[ G_t \nabla_{\vec{\psi}} \log \pi(\mathbf{a}_t|\boldsymbol{\theta}_t; \vec{\psi}) \right] \quad (\text{E.7})$$

1230 where  $G_t$  is the return at time  $t$ , and where we have written  $q_{\pi}(\boldsymbol{\theta}_t, \mathbf{a}_t) = \mathbb{E}_{\pi}[G_t|\boldsymbol{\theta}_t, \mathbf{a}_t]$ . The  
1231 argument of this expectation can be sampled on every timestep, and the expectation of this  
1232 quantity is equal to the gradient  $\nabla_{\vec{\psi}} J(\vec{\psi})$  [105]. This gives the update rule that we used in  
1233 Eq. 37:

$$\vec{\psi}_{t+1} = \vec{\gamma}_t + \alpha \mathbf{R}_t \nabla_{\vec{\psi}} \log \pi(\mathbf{a}_t|\boldsymbol{\theta}_t; \vec{\psi}) \quad (\text{E.8})$$

1234 where we replaced the full return  $G_t$  with the instantaneous reward  $R_t$  (i.e., we considered a  
1235 purely myopic agent).

## 1236 F Structuring a motor drive about a goal heading

1237 We consider a motor drive  $m$  of the form:

$$m = \int_0^{2\pi} d\theta \cos(\theta - \theta_c) f_g(\theta), \quad (\text{F.1})$$

1238 where  $\cos(\theta - \theta_c)$  is an activity profile that encodes the current compass heading  $\theta_c$ , and  
1239 where  $f_g(\theta)$  is an arbitrary activity profile that encodes the goal heading.

1240 Expanding the  $\cos$ , we have:

$$\begin{aligned} m &= \int_0^{2\pi} d\theta \frac{1}{2} (e^{i(\theta - \theta_c)} + e^{i(\theta - \theta_c)}) f_g(\theta) \\ &= \frac{e^{-i\theta_c}}{2} \int_0^{2\pi} d\theta e^{i\theta} f_g(\theta) + \frac{e^{i\theta_c}}{2} \int_0^{2\pi} d\theta e^{-i\theta} f_g(\theta), \end{aligned} \quad (\text{F.2})$$

1241 where

$$\int_0^{2\pi} d\theta e^{\pm i\theta} f_g(\theta) \equiv r_g e^{\pm i\theta_g} \quad (\text{F.3})$$

1242 defines the circular mean of the function  $f_g(\theta)$ , and  $r_g$  and  $\theta_g$  are the modulus and orientation  
1243 of the circular mean, respectively.

1244 This allows us to write:

$$\begin{aligned} m &= m(\theta_c, \theta_g) = \frac{r_g}{2} (e^{-i\theta_c} e^{i\theta_g} + e^{i\theta_c} e^{-i\theta_g}) \\ &= r_g \cos(\theta_c - \theta_g). \end{aligned} \quad (\text{F.4})$$

1245 Thus, when integrated over a full period, the product of the goal and compass activity profiles  
1246 will be largest when the compass heading is aligned with the circular mean of the goal activity  
1247 profile. We can then use this circular mean to define the goal heading.

1248 If we now add a phase shift  $\phi$  to the compass heading, this becomes:

$$\begin{aligned} m(\theta_c, \theta_g, \phi) &= \int_0^{2\pi} d\theta \cos(\theta - (\theta_c + \phi)) f_g(\theta) \\ &= r_g \cos(\theta_c + \phi - \theta_g), \end{aligned} \quad (\text{F.5})$$

1249 which peaks at  $\theta_c = \theta_g - \phi$ . Different phase shifts can then be used to generate motor drives  
1250 that peak at different compass headings relative to the goal heading. Importantly, the motor  
1251 drive will retain a sinusoidal profile regardless of the shape of  $f_g(\theta)$ , and thus regardless of  
1252 the goal heading  $\theta_g$ . As a result, we can use any method we like for updating the goal activity  
1253 profile, and it will not disrupt the structure of the motor drive.



## References

- 1254
- 1255 [1] B. Julesz, E. N. Gilbert and J. D. Victor, *Visual discrimination of textures with identical*  
1256 *third-order statistics*, Biological cybernetics **31**, 137 (1978).
- 1257 [2] J. D. Victor, C. Chubb and M. M. Conte, *Interaction of luminance and higher-order statis-*  
1258 *tics in texture discrimination*, Vision research **45**(3), 311 (2005).
- 1259 [3] G. Tkačik, J. S. Prentice, J. D. Victor and V. Balasubramanian, *Local statistics in natural*  
1260 *scenes predict the saliency of synthetic textures*, Proceedings of the National Academy of  
1261 Sciences **107**(42), 18149 (2010).
- 1262 [4] A. M. Hermundstad, J. J. Briguglio, M. M. Conte, J. D. Victor, V. Balasubramanian and  
1263 G. Tkačik, *Variance predicts salience in central sensory processing*, eLife **3**, e03722 (2014).
- 1264 [5] D. J. MacKay, *Information theory, inference and learning algorithms*, Cambridge univer-  
1265 sity press (2003).
- 1266 [6] E. P. Simoncelli and B. A. Olshausen, *Natural image statistics and neural representation*,  
1267 Annual review of neuroscience **24**(1), 1193 (2001).
- 1268 [7] F. Attneave, *Some informational aspects of visual perception.*, Psychological review **61**(3),  
1269 183 (1954).
- 1270 [8] H. B. Barlow *et al.*, *Possible principles underlying the transformation of sensory messages*,  
1271 Sensory communication **1**, 217 (1961).
- 1272 [9] S. Laughlin, *A simple coding procedure enhances a neuron's information capacity*,  
1273 Zeitschrift für Naturforschung c **36**(9-10), 910 (1981).
- 1274 [10] M. V. Srinivasan, S. B. Laughlin and A. Dubs, *Predictive coding: a fresh view of inhibition*  
1275 *in the retina*, Proceedings of the Royal Society of London. Series B. Biological Sciences  
1276 **216**(1205), 427 (1982).
- 1277 [11] J. J. Atick and A. N. Redlich, *Towards a theory of early visual processing*, Neural compu-  
1278 tation **2**(3), 308 (1990).
- 1279 [12] J. J. Atick and A. N. Redlich, *What does the retina know about natural scenes?*, Neural  
1280 computation **4**(2), 196 (1992).
- 1281 [13] J. H. van Hateren, *A theory of maximizing sensory information*, Biological cybernetics  
1282 **68**(1), 23 (1992).
- 1283 [14] G. Tkačik, J. S. Prentice, V. Balasubramanian and E. Schneidman, *Optimal popula-*  
1284 *tion coding by noisy spiking neurons*, Proceedings of the National Academy of Sciences  
1285 **107**(32), 14419 (2010).
- 1286 [15] E. Doi and M. S. Lewicki, *A simple model of optimal population coding for sensory systems*,  
1287 PLoS computational biology **10**(8), e1003761 (2014).
- 1288 [16] N. Brenner, W. Bialek and R. d. R. Van Steveninck, *Adaptive rescaling maximizes infor-*  
1289 *mation transmission*, Neuron **26**(3), 695 (2000).
- 1290 [17] A. L. Fairhall, G. D. Lewen, W. Bialek and R. R. d. R. van Steveninck, *Efficiency and*  
1291 *ambiguity in an adaptive neural code*, Nature **412**(6849), 787 (2001).

- 1292 [18] W. F. Młynarski and A. M. Hermundstad, *Efficient and adaptive sensory codes*, Nature  
1293 Neuroscience **24**(7), 998 (2021).
- 1294 [19] W. F. Młynarski and A. M. Hermundstad, *Adaptive coding for dynamic sensory inference*,  
1295 Elife **7**, e32055 (2018).
- 1296 [20] I. Dean, N. S. Harper and D. McAlpine, *Neural population coding of sound level adapts  
1297 to stimulus statistics*, Nature neuroscience **8**(12), 1684 (2005).
- 1298 [21] B. Wark, A. Fairhall and F. Rieke, *Timescales of inference in visual adaptation*, Neuron  
1299 **61**(5), 750 (2009).
- 1300 [22] D. B. Kastner and S. A. Baccus, *Coordinated dynamic encoding in the retina using opposing  
1301 forms of plasticity*, Nature neuroscience **14**(10), 1317 (2011).
- 1302 [23] I. Nemenman, *4. information theory and adaptation*, Quantitative Biology: From Molec-  
1303 ular to Cellular Systems p. 73 (2012).
- 1304 [24] C. Robert, *The Bayesian choice: from decision-theoretic foundations to computational  
1305 implementation*, Springer Science & Business Media (2007).
- 1306 [25] D. B. Kastner and S. A. Baccus, *Spatial segregation of adaptation and predictive sensiti-  
1307 zation in retinal ganglion cells*, Neuron **79**(3), 541 (2013).
- 1308 [26] T. R. Appleby and M. B. Manookin, *Neural sensitization improves encoding fidelity in the  
1309 primate retina*, Nature communications **10**(4017), 1 (2019).
- 1310 [27] B. N. Lundstrom, M. H. Higgs, W. J. Spain and A. L. Fairhall, *Fractional differentiation  
1311 by neocortical pyramidal neurons*, Nature neuroscience **11**(11), 1335 (2008).
- 1312 [28] S. E. Clarke, A. Longtin and L. Maler, *Contrast coding in the electrosensory system: par-  
1313 allels with visual computation*, Nature Reviews Neuroscience **16**(12), 733 (2015).
- 1314 [29] J. O'Doherty, M. L. Kringelbach, E. T. Rolls, J. Hornak and C. Andrews, *Abstract reward  
1315 and punishment representations in the human orbitofrontal cortex*, Nature neuroscience  
1316 **4**(1), 95 (2001).
- 1317 [30] P. Vertech, E. Lottem, D. Sarra, B. Godinho, I. Treves, T. Quendera, M. N. Oude  
1318 Lohuis and Z. F. Mainen, *Inference-based decisions in a hidden state foraging task:  
1319 Differential contributions of prefrontal cortical areas*, Neuron **106**(1), 166 (2020),  
1320 doi:<https://doi.org/10.1016/j.neuron.2020.01.017>.
- 1321 [31] L. P. Sugrue, G. S. Corrado and W. T. Newsome, *Matching behavior and the representation  
1322 of value in the parietal cortex*, Science **304**(5678), 1782 (2004).
- 1323 [32] V. D. Costa, V. L. Tran, J. Turchi and B. B. Averbeck, *Reversal learning and dopamine: a  
1324 bayesian perspective*, Journal of Neuroscience **35**(6), 2407 (2015).
- 1325 [33] B. A. Bari, C. D. Grossman, E. E. Lubin, A. E. Rajagopalan, J. I. Cressy and J. Y. Cohen,  
1326 *Stable representations of decision variables for flexible behavior*, Neuron **103**(5), 922  
1327 (2019), doi:<https://doi.org/10.1016/j.neuron.2019.06.001>.
- 1328 [34] M. P. Karlsson, D. G. Tervo and A. Y. Karpova, *Network resets in medial prefrontal cortex  
1329 mark the onset of behavioral uncertainty*, Science **338**(6103), 135 (2012).
- 1330 [35] Y. Liu, Y. Xin and N.-l. Xu, *A cortical circuit mechanism for structural knowledge-based  
1331 flexible sensorimotor decision-making*, Neuron **109**(12), 2009 (2021).

- 1332 [36] C. C. Beron, S. Q. Neufeld, S. W. Linderman and B. L. Sabatini, *Mice exhibit stochastic*  
1333 *and efficient action switching during probabilistic decision making*, Proceedings of the  
1334 National Academy of Sciences **119**(15), e2113961119 (2022).
- 1335 [37] A. Rajagopalan, R. Darshan, J. E. Fitzgerald and G. C. Turner, *Expectation-*  
1336 *based learning rules underlie dynamic foraging in drosophila*, bioRxiv (2022),  
1337 doi:[10.1101/2022.05.24.493252](https://doi.org/10.1101/2022.05.24.493252).
- 1338 [38] K. J. Astrom, *Optimal control of markov decision processes with incomplete state estima-*  
1339 *tion*, J. Math. Anal. Applic. **10**, 174 (1965).
- 1340 [39] L. P. Kaelbling, M. L. Littman and A. R. Cassandra, *Planning and acting in partially*  
1341 *observable stochastic domains*, Artificial intelligence **101**(1-2), 99 (1998).
- 1342 [40] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, The MIT Press,  
1343 second edn. (2018).
- 1344 [41] T. Ma and A. M. Hermundstad, *A vast space of compact strategies for highly efficient*  
1345 *decisions*, bioRxiv pp. 2022–08 (2022).
- 1346 [42] N. Tishby and D. Polani, *Information theory of decisions and actions*, In *Perception-action*  
1347 *cycle: Models, architectures, and hardware*, pp. 601–636. Springer (2010).
- 1348 [43] H. S. Seung, *How the brain keeps the eyes still*, Proceedings of the National Academy of  
1349 Sciences **93**(23), 13339 (1996), doi:[10.1073/pnas.93.23.13339](https://doi.org/10.1073/pnas.93.23.13339).
- 1350 [44] M. Camperi and X.-J. Wang, *A model of visuospatial working memory in prefrontal cortex:*  
1351 *recurrent network and cellular bistability*, Journal of computational neuroscience **5**(4),  
1352 383 (1998).
- 1353 [45] D. Hansel and H. Sompolinsky, *Modeling feature selectivity in local cortical circuits*, In  
1354 C. Koch and I. Segev, eds., *Methods in neuronal modeling: from synapses to networks*,  
1355 chap. Chap 13. Cambridge, MA, 2 edn. (1998).
- 1356 [46] M. Tsodyks, *Attractor neural network models of spatial maps in hippocampus*, Hippocam-  
1357 pus **9**(4), 481 (1999).
- 1358 [47] A. Compte, N. Brunel, P. S. Goldman-Rakic and X.-J. Wang, *Synaptic mechanisms and*  
1359 *network dynamics underlying spatial working memory in a cortical network model*, Cere-  
1360 bral cortex **10**(9), 910 (2000).
- 1361 [48] A. Samsonovich and B. L. McNaughton, *Path integration and cognitive mapping in a con-*  
1362 *tinuous attractor neural network model*, Journal of Neuroscience **17**(15), 5900 (1997).
- 1363 [49] S. Stringer, T. Trappenberg, E. Rolls and I. de Araujo, *Self-organizing continuous attractor*  
1364 *networks and path integration: one-dimensional models of head direction cells.*, Network  
1365 (Bristol, England) **13**(2), 217 (2002).
- 1366 [50] Y. Burak and I. R. Fiete, *Accurate path integration in continuous attractor network models*  
1367 *of grid cells*, PLoS computational biology **5**(2), e1000291 (2009).
- 1368 [51] H. S. Seung, *Continuous attractors and oculomotor control*, Neural Networks **11**(7-8),  
1369 1253 (1998).
- 1370 [52] M. Goldman, A. Compte and X.-J. Wang, *Neural integrator models*, In L. R. Squire, ed.,  
1371 *Encyclopedia of Neuroscience*, pp. 165–178. Academic Press, Oxford, ISBN 978-0-08-  
1372 045046-9, doi:<https://doi.org/10.1016/B978-008045046-9.01434-0> (2009).

- 1373 [53] J. J. Knierim and K. Zhang, *Attractor dynamics of spatially correlated neural ac-*  
1374 *tivity in the limbic system*, Annual Review of Neuroscience **35**(1), 267 (2012),  
1375 doi:[10.1146/annurev-neuro-062111-150351](https://doi.org/10.1146/annurev-neuro-062111-150351), PMID: 22462545.
- 1376 [54] R. Chaudhuri and I. Fiete, *Computational principles of memory*, Nature neuroscience  
1377 **19**(3), 394 (2016).
- 1378 [55] B. K. Hulse and V. Jayaraman, *Mechanisms underlying the neural computation of head*  
1379 *direction*, Annual review of neuroscience **43**, 31 (2020).
- 1380 [56] K. Wimmer, D. Q. Nykamp, C. Constantinidis and A. Compte, *Bump attractor dynamics*  
1381 *in prefrontal cortex explains behavioral precision in spatial working memory*, Nature  
1382 neuroscience **17**(3), 431 (2014).
- 1383 [57] I. Pisokas, S. Heinze and B. Webb, *The head direction circuit of two insect species*, eLife  
1384 **9**, e53985 (2020), doi:[10.7554/eLife.53985](https://doi.org/10.7554/eLife.53985).
- 1385 [58] D. B. Turner-Evans, S. Wegener, H. Rouault, R. Franconville, T. Wolff, J. D. Seelig,  
1386 S. Druckmann and V. Jayaraman, *Angular velocity integration in a fly heading circuit*,  
1387 eLife **6**, e23496 (2017), doi:<https://doi.org/10.7554/eLife.23496>.
- 1388 [59] K. S. Kakaria and B. L. de Bivort, *Ring attractor dynamics emerge from a spiking model*  
1389 *of the entire protocerebral bridge*, Frontiers in Behavioral Neuroscience **11** (2017),  
1390 doi:[10.3389/fnbeh.2017.00008](https://doi.org/10.3389/fnbeh.2017.00008).
- 1391 [60] J. D. Seelig and V. Jayaraman, *Neural dynamics for landmark orientation and angular*  
1392 *path integration*, Nature **521**(7551), 186 (2015), doi:[10.1038/nature14446](https://doi.org/10.1038/nature14446).
- 1393 [61] J. Green, A. Adachi, K. K. Shaw, J. D. Hirokawa, P. S. Magani and G. Maimon, *A neu-*  
1394 *ral circuit architecture for angular integration in drosophila*, Nature **546**, 101 (2017),  
1395 doi:<https://doi.org/10.1038/nature22343>.
- 1396 [62] D. B. Turner-Evans, K. T. Jensen, S. Ali, T. Paterson, A. Sheridan, R. P. Ray, T. Wolff,  
1397 J. S. Lauritzen, G. M. Rubin, D. D. Bock and V. Jayaraman, *The neuroanatomical ul-*  
1398 *trastructure and function of a biological ring attractor*, Neuron **108**(1), 145 (2020),  
1399 doi:<https://doi.org/10.1016/j.neuron.2020.08.006>.
- 1400 [63] B. K. Hulse, H. Haberkern, R. Franconville, D. B. Turner-Evans, S.-y. Takemura, T. Wolff,  
1401 M. Noorman, M. Dreher, C. Dan, R. Parekh *et al.*, *A connectome of the drosophila cen-*  
1402 *tral complex reveals network motifs suitable for flexible navigation and context-dependent*  
1403 *action selection*, ELife **10**, e66039 (2021).
- 1404 [64] S. S. Kim, H. Rouault, S. Druckmann and V. Jayaraman, *Ring attractor*  
1405 *dynamics in the drosophila central brain*, Science **356**(6340), 849 (2017),  
1406 doi:[10.1126/science.aal4835](https://doi.org/10.1126/science.aal4835).
- 1407 [65] M. Noorman, B. K. Hulse, V. Jayaraman, S. Romani and A. M. Hermundstad, *Accurate*  
1408 *angular integration with only a handful of neurons*, bioRxiv pp. 2022–05 (2022).
- 1409 [66] S. H. Strogatz, *Nonlinear dynamics and chaos with student solutions manual: With ap-*  
1410 *plications to physics, biology, chemistry, and engineering*, CRC press (2018).
- 1411 [67] B. K. Hulse, A. Stanoev, D. B. Turner-Evans, J. Seelig and V. Jayaraman, *A rotational*  
1412 *velocity estimate constructed through visuomotor competition updates the fly's neural com-*  
1413 *pass*, bioRxiv pp. 2023–09 (2023).

- 1414 [68] J. Laurens and D. E. Angelaki, *The brain compass: a perspective on how self-motion*  
1415 *updates the head direction cell attractor*, *Neuron* **97**(2), 275 (2018).
- 1416 [69] K. Zhang, *Representation of spatial orientation by the intrinsic dynamics of the head-*  
1417 *direction cell ensemble: a theory*, *Journal of Neuroscience* **16**(6), 2112 (1996),  
1418 doi:[10.1523/JNEUROSCI.16-06-02112.1996](https://doi.org/10.1523/JNEUROSCI.16-06-02112.1996).
- 1419 [70] H. B.j., J. Omoto, P. Kandimalla, B. Nguyen, M. Keleş, N. Boyd, V. Hartenstein and  
1420 M. Frye, *A visual pathway for skylight polarization processing in drosophila*, *eLife* **10**,  
1421 e63225 (2021), doi:[10.7554/eLife.63225](https://doi.org/10.7554/eLife.63225).
- 1422 [71] S. Heinze and S. M. Reppert, *Sun compass integration of skylight*  
1423 *cues in migratory monarch butterflies*, *Neuron* **69**(2), 345 (2011),  
1424 doi:<https://doi.org/10.1016/j.neuron.2010.12.025>.
- 1425 [72] H. Vitzthum, M. Müller and U. Homberg, *Neurons of the central complex of the locust*  
1426 *schistocerca gregaria are sensitive to polarized light*, *Journal of Neuroscience* **22**(3), 1114  
1427 (2002), doi:[10.1523/JNEUROSCI.22-03-01114.2002](https://doi.org/10.1523/JNEUROSCI.22-03-01114.2002).
- 1428 [73] T. Okubo, P. Patella, I. D'Alessandro and R. Wilson, *A neural network for wind-guided*  
1429 *compass navigation*, *Neuron* **107**, 924 (2020), doi:[10.1016/j.neuron.2020.06.022](https://doi.org/10.1016/j.neuron.2020.06.022).
- 1430 [74] J. Seelig and V. Jayaraman, *Feature detection and orientation tuning in the drosophila cen-*  
1431 *tral complex*, *Nature* **503**, 262–266 (2013), doi:<https://doi.org/10.1038/nature12601>.
- 1432 [75] J. J. Omoto, M. F. Keleş, B.-C. M. Nguyen, C. Bolanos, J. K. Lovick, M. A. Frye  
1433 and V. Hartenstein, *Visual input to the drosophila central complex by developmentally*  
1434 *and functionally distinct neuronal populations*, *Current Biology* **27**(8), 1098 (2017),  
1435 doi:<https://doi.org/10.1016/j.cub.2017.02.063>.
- 1436 [76] Y. Sun, A. Nern, R. Franconville, H. Dana, E. R. Schreiter, L. L. Looger, K. Svo-  
1437 boda, D. S. Kim, A. M. Hermundstad and V. Jayaraman, *Neural signatures of*  
1438 *dynamic stimulus selection in drosophila*, *Nat Neurosci* **20**, 1104–1113 (2017),  
1439 doi:<https://doi.org/10.1038/nn.4581>.
- 1440 [77] H. Shiozaki and H. Kazama, *Parallel encoding of recent visual experience and self-*  
1441 *motion during navigation in drosophila*, *Nat Neurosci* **20**, 1395–1403 (2017),  
1442 doi:<https://doi.org/10.1038/nn.4628>.
- 1443 [78] D. W. Dong and J. J. Atick, *Temporal decorrelation: a theory of lagged and nonlagged*  
1444 *responses in the lateral geniculate nucleus*, *Network: Computation in neural systems*  
1445 **6**(2), 159 (1995).
- 1446 [79] B. A. Olshausen and D. J. Field, *Emergence of simple-cell receptive field properties by*  
1447 *learning a sparse code for natural images*, *Nature* **381**(6583), 607 (1996).
- 1448 [80] A. J. Cope, C. Sabo, E. Vasilaki, A. B. Barron and J. A. R. Marshall, *A computational*  
1449 *model of the integration of landmarks and motion in the insect central complex*, *PLOS*  
1450 *ONE* **12**, 1 (2017), doi:[10.1371/journal.pone.0172325](https://doi.org/10.1371/journal.pone.0172325).
- 1451 [81] J. Green and G. Maimon, *Building a heading signal from anatomically defined neu-*  
1452 *ron types in the drosophila central complex*, *Current Opinion in Neurobiology* **52**, 156  
1453 (2018), doi:<https://doi.org/10.1016/j.conb.2018.06.010>, *Systems Neuroscience*.



- 1454 [82] S. S. Kim, A. M. Hermundstad, S. Romani, L. F. Abbott and V. Jayaraman, *Generation of*  
1455 *stable heading representations in diverse visual scenes*, *Nature* **576**(7785), 126 (2019),  
1456 doi:[10.1038/s41586-019-1767-1](https://doi.org/10.1038/s41586-019-1767-1).
- 1457 [83] Y. E. Fisher, J. Lu, I. D'Alessandro and R. I. Wilson, *Sensorimotor experience*  
1458 *remaps visual input to a heading-direction network*, *Nature* **576**(7785), 121 (2019),  
1459 doi:[10.1038/s41586-019-1772-4](https://doi.org/10.1038/s41586-019-1772-4).
- 1460 [84] Y. E. Fisher, M. Marquis, I. D'Alessandro and R. I. Wilson, *Dopamine promotes head*  
1461 *direction plasticity during orienting movements*, *Nature* **612**(7939), 316 (2022).
- 1462 [85] H. Haberkern, S. S. Chitnis, P. M. Hubbard, T. Goulet, A. M. Hermundstad and V. Ja-  
1463 yaraman, *Maintaining a stable head direction representation in naturalistic visual envi-*  
1464 *ronments*, *bioRxiv* (2022), doi:[10.1101/2022.05.17.492284](https://doi.org/10.1101/2022.05.17.492284).
- 1465 [86] T. Kohonen, *The self-organizing map*, *Proceedings of the IEEE* **78**(9), 1464 (1990).
- 1466 [87] U. Hanesch, K.-F. Fischbach and M. Heisenberg, *Neuronal architecture of the central*  
1467 *complex in drosophila melanogaster*, *Cell and Tissue Research* **257**(2), 343 (1989),  
1468 doi:[10.1007/BF00261838](https://doi.org/10.1007/BF00261838).
- 1469 [88] J. Isaacman-Beck, K. C. Paik, C. F. R. Wienecke, H. H. Yang, Y. E. Fisher, I. E. Wang,  
1470 I. G. Ishida, G. Maimon, R. I. Wilson and T. R. Clandinin, *Sparc enables genetic ma-*  
1471 *nipulation of precise proportions of cells*, *Nature Neuroscience* **23**(9), 1168 (2020),  
1472 doi:[10.1038/s41593-020-0668-9](https://doi.org/10.1038/s41593-020-0668-9).
- 1473 [89] R. Wolf and M. Heisenberg, *Basic organization of operant behavior as revealed in*  
1474 *drosophila flight orientation*, *Journal of Comparative Physiology A* **169**(6), 699 (1991),  
1475 doi:[10.1007/BF00194898](https://doi.org/10.1007/BF00194898).
- 1476 [90] P. W. Tillberg, personal communication.
- 1477 [91] K. G. Götz, *Course-control, metabolism and wing interference during ultralong tethered*  
1478 *flight in drosophila melanogaster*, *Journal of Experimental Biology* **128**(1), 35 (1987),  
1479 doi:[10.1242/jeb.128.1.35](https://doi.org/10.1242/jeb.128.1.35).
- 1480 [92] M. B. Reiser and M. H. Dickinson, *A modular display system for insect be-*  
1481 *havioral neuroscience*, *Journal of Neuroscience Methods* **167**(2), 127 (2008),  
1482 doi:<https://doi.org/10.1016/j.jneumeth.2007.07.019>.
- 1483 [93] C. Dan, B. K. Hulse, V. Jayaraman and A. M. Hermundstad, *A neural circuit architecture*  
1484 *for rapid behavioral flexibility in goal-directed navigation*, *bioRxiv* (2024).
- 1485 [94] R. Wolf and M. Heisenberg, *On the fine structure of yaw torque in visual flight*  
1486 *orientation of drosophila melanogaster*, *J. Comp. Physiol. A* **140**, 69–80 (1980),  
1487 doi:<https://doi.org/10.1007/BF00613749>.
- 1488 [95] F. T. Muijres, M. J. Elzinga, N. A. Iwasaki and M. H. Dickinson, *Body saccades of*  
1489 *drosophila consist of stereotyped banked turns*, *Journal of Experimental Biology* **218**,  
1490 864 (2015), doi:[doi: 10.1242/jeb.114280](https://doi.org/10.1242/jeb.114280).
- 1491 [96] E. Schrödinger, *Zur theorie der fall-und steigversuche an teilchen mit brown'scher bewe-*  
1492 *gung*, *Physikalische Zeitschrift* **16**, 289 (1915).
- 1493 [97] M. Smoluchowski, *Notiz über die berechnung der brown'schen molekularbewegung bei der*  
1494 *ehrenhaft-millikan'schen versuchsanordnung*, *Pisma Mariana Smoluchowskiego* **1**(2),  
1495 477 (1927).



- 1496 [98] J. L. Folks and R. S. Chhikara, *The inverse gaussian distribution and its statistical applica-*  
1497 *tion—a review*, Journal of the Royal Statistical Society Series B: Statistical Methodology  
1498 **40**(3), 263 (1978).
- 1499 [99] A. Y. Ng, S. Russell *et al.*, *Algorithms for inverse reinforcement learning.*, In *Icml*, vol. 1,  
1500 p. 2 (2000).
- 1501 [100] D. S. Touretzky, A. D. Redish and H. S. Wan, *Neural representation of space Using sinu-*  
1502 *oidal arrays*, Neural Computation **5**(6), 869 (1993), doi:[10.1162/neco.1993.5.6.869](https://doi.org/10.1162/neco.1993.5.6.869).
- 1503 [101] R. Goulard, C. Buehlmann, J. E. Niven, P. Graham and B. Webb, *A unified mechanism*  
1504 *for innate and learned visual landmark guidance in the insect central complex*, PLoS  
1505 computational biology **17**(9), e1009383 (2021).
- 1506 [102] C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer-Verlag, second edn.  
1507 (2006).
- 1508 [103] M. Botvinick, S. Ritter, J. X. Wang, Z. Kurth-Nelson, C. Blundell and D. Hassabis,  
1509 *Reinforcement learning, fast and slow*, Trends in Cognitive Sciences **23**(5) (2019),  
1510 doi:<https://doi.org/10.1016/j.tics.2019.02.006>.
- 1511 [104] R. E. Bellman, *Dynamic Programming*, Princeton University Press (1957).
- 1512 [105] R. J. Williams, *Simple statistical gradient-following algorithms for connectionist reinfor-*  
1513 *ment learning*, Machine learning **8**, 229 (1992).