Chapter 6: Attention

The world is filled with stimuli, but only some of them are important at any given time. The brain focuses on important stimuli through attentional mechanisms. This chapter considers two mechanisms. One reflects expectations based on prior probabilities: stimuli that are expected will tend to receive higher posterior probability, rendering them more detectable. A second mechanism reflects resource allocation: attentional cues direct limited resources to important stimuli, amplifying information from those signals.

Attention is the enhanced processing of certain signals, typically (but not always) at the expense of other signals. It can be induced "top-down" (e.g., by preparatory cues, volitional orienting) or "bottom-up" (e.g., capture of attention by salient or unexpected stimuli). Much of what we know about attention concerns sensory (particularly visual) signals, and thus this chapter will focus on visual attention, but we note here that attention is a more general phenomenon that can be applied internally to cognitive processes such as memory, imagery, and self-directed thought (Chun et al., 2011).

1 *Attention as prior probability*

The most straightforward way to think about attention in a Bayesian framework is through the prior: if some events have higher prior probability, they will be more easily detected.

1.1 *Cued attention*

To make this concrete, we will use the classic Posner cueing task (Posner, 1980), illustrated in Figure 1. On each trial, a central cue appears, either a fixation cross or an arrow pointing left/right. Subsequently a target appears either on the left or the right, and the subject is asked to press a key as quickly as possible when they've detected the target. The trial types differ only in the position uncertainty implied by the cue. Targets appear on the left or the right with equal frequency following the fixation cross; thus position uncertainty is 50%. Targets appear in the direction indicated by the arrow with probability 80%. This produces both valid trials (when the target appeared in the expected location) and invalid trials (when the target appeared in the unexpected location). The main behavioral finding is that response time is longest on invalid trials and shortest on valid trials (the difference between these two response times is known as the *validity effect*). Intuitively, the cue orients attention to the expected spatial location.

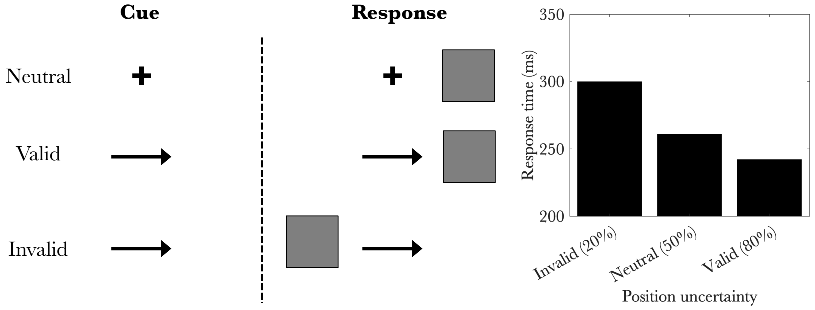

Figure 1: **The Posner cueing task**. Data replotted from Posner (1980).

An important aspect of the Posner task is that subjects are asked to maintain central fixation throughout the trial; eye position is continuously monitored, and trials are rejected if fixation is broken. This means that the orienting of attention is "covert" in the sense that the sensory signals are largely unchanged. Only the internal direction of attention is altered by the cues.

We can now start to develop a model of performance in the Posner task. The hidden state is binary, $s \in \{\text{Left}, \text{Right}\}$, and is encoded by neural activity x at the time of the target appearance. As in earlier chapters, we will assume that the spike count of neuron d is Poisson-distributed with rate $f_d(s)$. Recall from Chapter 4 how evidence accumulation can be implemented by a readout neuron whose input current $I(t)$ linear weights spikes from the encoding population:

$$I(t) = \sum_d w_d z_d(t), \qquad (1)$$

where $z_d(t) = 1$ if neuron d spikes at time t (0 otherwise), and the synaptic strength is given by:

$$w_d = \log \frac{f_d(\text{Left})}{f_d(\text{Right})}. \qquad (2)$$

The input current is integrated by the membrane potential $\mu(t)$:

$$C\dot{\mu} = I(t), \qquad \mu(0) = \mu^0, \qquad (3)$$

where C is the membrane capacitance and μ^0 is the potential at the time of cue presentation ($t = 0$). For present purposes, μ^0 is the critical variable, because that's what encodes the prior:

Importantly, the prior is cue-dependent. The cue itself is encoded by a separate population of neurons that send input to the readout neuron.

$$\mu^0 = \log \frac{p(\text{Left}|\text{Cue})}{p(\text{Right}|\text{Cue})}. \qquad (4)$$

With this setup, the membrane potential at time t represents the posterior log odds. We can link the model to behavior by assuming that a detection response is made whenever the firing rate of the readout neuron crosses one of two thresholds representing the two spatial locations.

See Chapter 7 for more on the decision-theoretic aspects of threshold-crossing.

The model can explain the main behavioral finding: the arrow cue drives the membrane potential up or down depending on its direction, which in turn speeds the crossing of the decision threshold when the target appears in the expected location. When it appears in the unexpected location, the evidence needs to move in the opposite direction, making response time longer—this formalizes the concept of "reorienting" that Posner and others have invoked to explain responses on invalid trials.

Monkeys can be taught to use symbolic predictive cues (Bowman et al., 1993), but their attention can also be directed by a peripheral cue flashed on one side of the screen. Neurons in parietal cortex exhibit elevated responding to targets that appear on the same side as the cue (Robinson et al., 1995), consistent with additional drive provided by the spatial prior. The physiology is also consistent with the observation that parietal lesions disrupt attentional orienting (Posner et al., 1984).

1.2 *Acetylcholine and cued attention*

A large literature implicates the neuromodulator acetylcholine in attention (Hasselmo and Sarter, 2011). Here we focus on the role of acetylcholine in cued attention, which received an influential theoretical treatment by Yu and Dayan (2005). Their account posited that acetylcholine reports *expected uncertainty* about target stimuli conditional on cues. Thus, acetylcholine should control the size of the validity effect: greater uncertainty about the cued target location will increase response times on valid trials and decrease response times on invalid trials.

In support of this model, Yu and Dayan cite work by Phillips et al. (2000), which manipulated acetylcholine levels pharmacologically in rodents performing the Posner task. Nicotine (an acetylcholine agonist) decreased the validity effect, while scopolamine (an acetylcholine antagonist) increased it. However, this study appears to be an outlier, because several other studies found the opposite effects: increased validity effect with nicotine, and decreased validity effect with scopolamine (Davidson et al., 1999) or another acetylcholine antagonist, mecamylamine (Stewart et al., 2001). The important take-away is that acetylcholine is clearly involved in cued attention, but perhaps not in the specific way posited by Yu and Dayan. The bulk

An agonist is a molecule that mimics the action of another molecule. An antagonist has the opposite effect, blocking action of the molecule.

of data are compatible with a model in which acetylcholine reports *expected certainty*.

One way to formalize this role in terms of the model described above would be to posit that acetylcholine reports μ^0. Thus, suppressing acetylcholine should have the effect of abolishing cue-dependent attentional modulation of the readout neuron. Consistent with this hypothesis, Davidson and Marrocco (2000) showed that attentional modulation of parietal cells during performance of the Posner task is suppressed in a dose-dependent manner by scopolamine.

Hasselmo and Sarter (2011) describe a possible biophysical mechanism for acetylcholine's attentional effects. When a neuron is sufficiently depolarized (by cue-evoked glutamate currents), calcium channels are opened, resulting in the activation of a calcium-sensitive nonspecific cation current, which causes further depolarization, further calcium influx, and so on—a self-sustaining loop producing persistent spiking. As a consequence, subsequent stimuli will be more effective at activating neurons with elevated activity.

A cation is a positively charged ion (e.g., sodium, calcium). A nonspecific cation current is a membrane current that does not depend on a particular ion.

1.3 *Effects on tuning*

Spatial priors can also be used to orient attention towards particular features. For example, McAdams and Maunsell (1999) studied orientation-tuned neurons in visual area V4 while monkeys performed a delayed match-to-sample task (Figure 2). On each trial, monkeys were presented with sample stimuli (a colored stimulus and an oriented grating). After a delay, they were presented with test stimuli (another colored stimulus and oriented grating) and instructed to judge whether one of the test stimuli matched the corresponding sample stimulus. Because each V4 neuron is also spatially tuned, and the gratings were always presented within the receptive field of the neuron, the experimenters could classify each grating as either "attended" (on orientation judgment trials) or "unattended" (on color judgment trials). The key finding was that attention multiplicatively enhanced the orientation tuning function.

One way to model this phenomenon is to assume that the prior is higher in the attended location (Dayan and Zemel, 1999). In the extreme case (zero prior probability outside the attended location), the resulting posterior takes the form of a "spotlight" that selectively amplifies information in a specific location. Rao (2005) developed a model of the delayed match-to-sample task in which the V4 neurons were modeled as reporting the posterior probability over orientation (s_1) given lower-level visual inputs (x), marginalizing over location

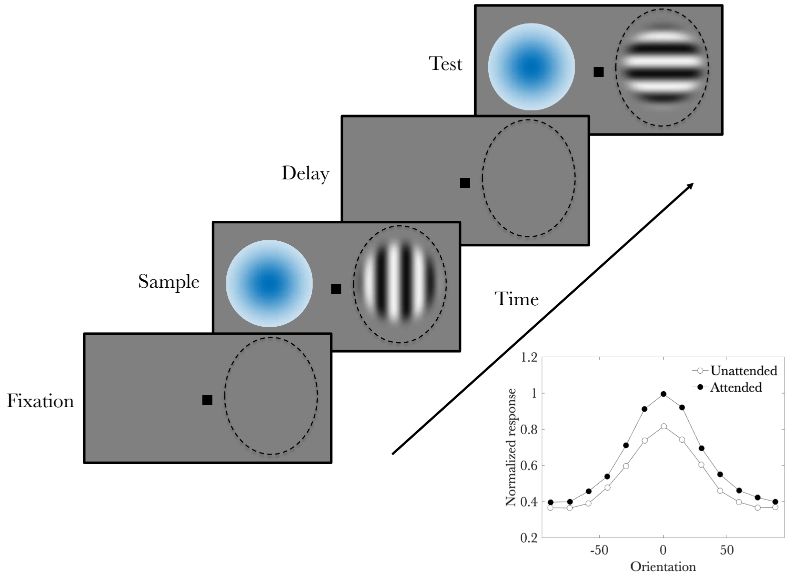

Figure 2: **Delayed match-to-sample task**. Responses are normalized to the peak firing rate in the attended condition. Dashed oval illustrates the receptive field of a V4 neuron. Replotted from McAdams and Maunsell (1999).

(s_2):

$$p(s_1|x) = \sum_{s_2} p(s_1|s_2, x)p(s_2|x), \tag{5}$$

where for simplicity we have again assumed a discrete set of spatial locations, $s_2 \in \{\text{Left}, \text{Right}\}$. Because $p(s_2|x)$ increases with $p(s_2)$, the posterior over orientation is modulated multiplicatively by the prior over location.

The same model can explain competitive attentional phenomena. When multiple stimuli are presented in the spatial receptive field of a V4 neuron, activity is suppressed (Reynolds et al., 1999). Figure 3 shows an example from recordings of a neuron tuned to vertical orientation when a nearby horizontal bar is simultaneously presented. The model explains this suppression due to uncertainty about orientation. When a monkey is trained to attend to the location where the horizontal bar is presented, activity is restored—a pattern recapitulated by the model due to the elevated prior probability assigned to the location.

1.4 Connection to normalization

The normalization model of attention (Reynolds and Heeger, 2009) has been used to compactly explain a very wide range of findings, and for this reason is one of the most influential modern accounts of attentional effects (particularly in visual cortex). Here we show that it is closely connected to the Bayesian account described above.

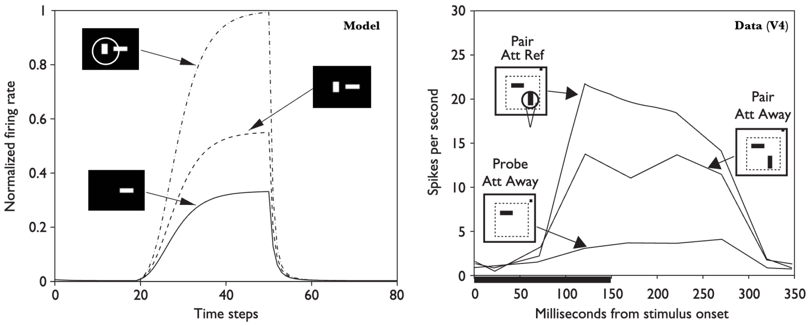

Figure 3: **Attentional competition in V4**. (Left) Model. (Right) Data. Adapted from Rao (2005); data from Reynolds et al. (1999).

The normalization model describes the firing rate of a neuron as a function of stimulus feature (s_1, typically orientation) and location (s_2), which we write as $s = (s_1, s_2)$:

$$f(s) = \frac{A(s)E(s)}{S + \nu}, \tag{6}$$

where ν is a gain control parameter, $E(s)$ is the stimulus drive, $A(s)$ is the attention field, and S is the suppressive drive derived from pooling the excitatory drive $E(s)A(s)$ over a wider range of s (the suppressive field).

For simplicity we ignore the spiking nonlinearity in the normalization equation.

We can map this onto the Bayesian model if we think of the likelihood $p(x|s)$ as encoding the stimulus drive, the prior $p(s)$ as encoding the attention field, and the marginal likelihood $p(x) = \sum_s p(x|s)p(s)$ as encoding the suppressive drive (pooled excitatory drive). Bayes' rule then implements a form of normalization with $\nu = 0$:

$$p(s|x) = \frac{p(x|s)p(s)}{p(x)}. \tag{7}$$

We can leverage this connection to elucidate how the Bayesian model explains some of the phenomena that originally motivated the normalization model.

Reynolds and Heeger (2009) noted a discrepancy between the results of several studies. When stimulus contrast is manipulated, it is possible to visualize the effect of attention on the relationship between stimulus contrast and firing rate (the contrast response function). For neurons in visual areas (e.g., V4), firing rate generally increases with stimulus contrast (Figure 4). Some studies (e.g., Reynolds et al., 2000) have observed a multiplicative modulation of the contrast response function by attention—a bigger boost in firing for higher-contrast attended stimuli (though this boost saturates at very high contrasts). Other studies (e.g., Williford and Maunsell, 2006) have observed a shift in the contrast response function for attended stimuli—roughly the same boost in firing for all contrast levels. To explain this discrepancy, Reynolds and Heeger (2009) pointed out that Williford and Maunsell (2006) used larger stimuli and task demands that required more focal spatial attention (guiding precise saccades to a target location), whereas Reynolds et al. (2000) used smaller stimuli without demanding precise spatial attention. Putting all this together, Reynolds and Heeger proposed that multiplicative modulation arises when the attention field is small relative to the stimulus, whereas a shift arises when the attention field is large relative to the stimulus.

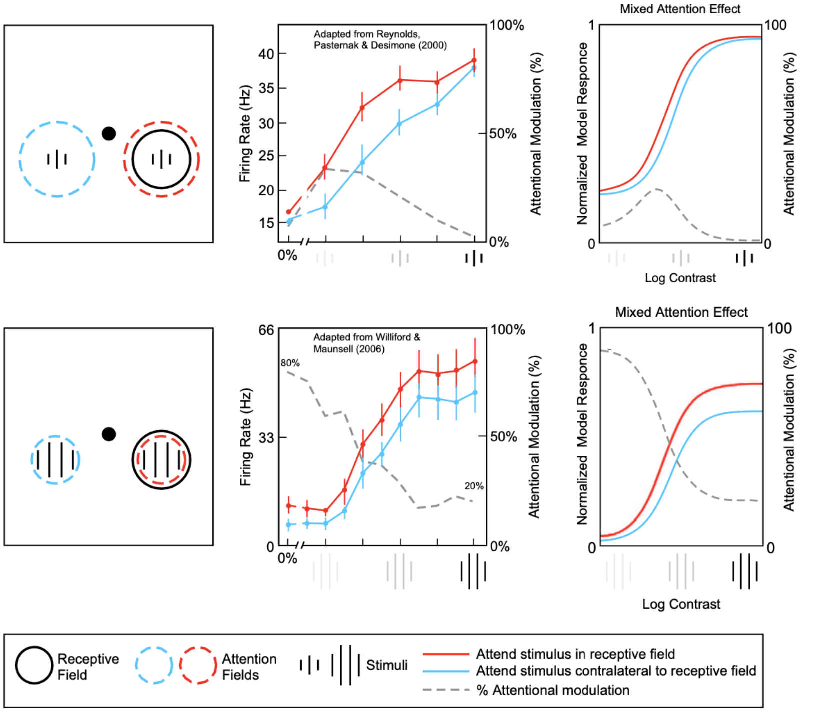

Figure 4: **Attentional modulation of contrast response in V4**. (Top) Small stimuli (Reynolds et al., 2000); (Bottom) Large stimuli (Williford and Maunsell, 2006). The left column shows a schematic of the experimental stimulus and the hypothetical attentional field. The middle column shows firing rate as a function of stimulus contrast. The right column shows model simulations. Adapted Reynolds and Heeger (2009).

To show how this explanation can be derived from the normalization model, Reynolds and Heeger assumed that stimulus drive is a linear function of contrast $g(s)$: $E(s) = \alpha g(s)$. They also assumed that the attentional field is approximately uniform with strength γ. In the case where the attention field is large relative to the stimulus, the attention field is approximately constant as a function of s and the contrast response function takes the following form:

$$f(s) \approx \frac{\alpha \gamma g(s)}{\gamma g(s) + \nu}. \tag{8}$$

This equation makes clear that attention multiplicatively interacts with contrast, as in Reynolds et al. (2000). In the case where the attention field is small relative to the stimulus, the attention field is approximately a focal point at the stimulus location, and the contrast response function takes the following form:

$$f(s) \approx \frac{\alpha\gamma g(s)}{\gamma g(s) + \omega g(s) + \nu}, \tag{9}$$

where $\omega \in [0,1]$ scales the suppressive drive from the region surrounding the focal point; ω increases with stimulus size. This equation predicts attention modulation across the range of contrasts (though not perfectly uniform across the range). To see this, note that when $g(s) \ll \nu$,

$$f(s) \approx \frac{\alpha\gamma g(s)}{\nu}, \tag{10}$$

which is increasing in γ. When $g(s) \gg \nu$,

$$f(s) \approx \frac{\alpha\gamma}{\gamma + \omega}, \tag{11}$$

which is now independent of $g(s)$ but nonetheless increasing in γ. Reynolds and Heeger showed that this regime can give rise to an attentional shift in the contrast response function similar to what was observed experimentally by Williford and Maunsell (2006).

The Bayesian model provides a normative interpretation of these ideas. Posterior probability increases with contrast due to its action on the likelihood (expressing the stimulus drive). Posterior probability also increase with attention due to its action on the prior, as discussed above. The marginal likelihood expresses the suppressive drive, increasing with stimulus size due to the fact that more stimulus locations enter into the normalizing constant. As shown in Figure 5, a Bayesian model formalizing this interpretation (Chikkerur et al., 2010) can capture the diverse attentional effects observed experimentally.

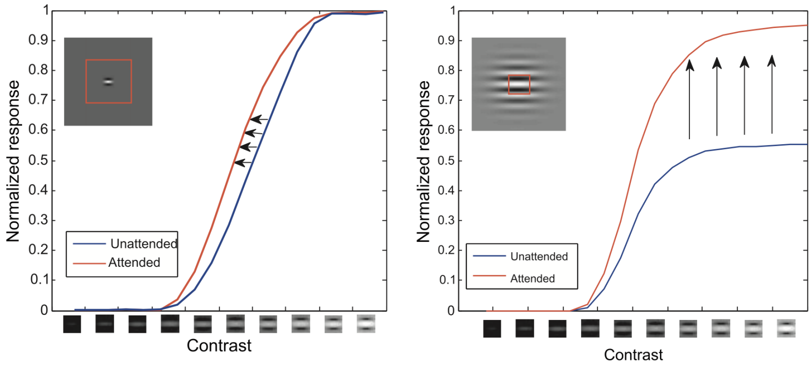

Figure 5: **Stimulus size determines the relationship between stimulus contrast and attention**. (Left) When the stimulus is small relative to the attention field, the contrast response function is shifted by attention. (Right) When the stimulus is large relative to the attention field, the contrast response function is scaled multiplicatively by attention. Reproduced from Chikkerur et al. (2010).

The normalization model can also explain why attention sometimes sharpens stimulus tuning (Spitzer et al., 1988; Martinez-Trujillo and Treue, 2004). Note that this contrasts with the result from McAdams and Maunsell (1999), where the entire tuning curve is shifted upward. A key factor responsible for this discrepancy is the variable over which attention is specified in different studies. In the McAdams and Maunsell study, attention was specified spatially (i.e., whether the stimulus was inside or outside the neuron's spatial receptive field). In the Martinez-Trujillo and Treue study, attention was specified based on the stimulus feature (motion direction in this case). This corresponds, in the normalization model, to amplifying the attention field near the attended feature value, with the effect that tuning is sharpened at those feature values (Figure 6). In Bayesian terms, this reflects the increase in prior probability for particular feature values; the qualitatively different effects on tuning functions in the two experimental paradigms are reduced to differences in spatial vs. feature-based priors.

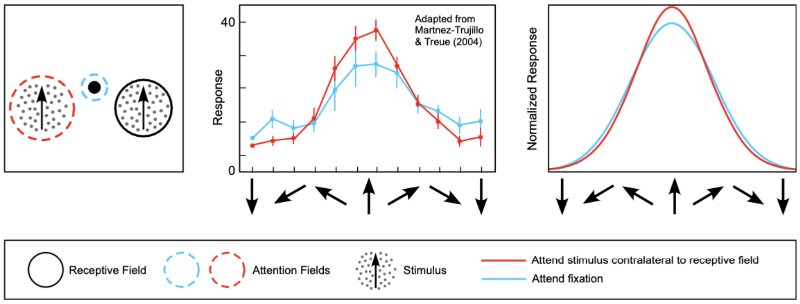

Figure 6: **Attention can sharpen tuning functions.** (Left) Experimental setup from Martinez-Trujillo and Treue (2004). One stimulus was always shown in the receptive field of the recorded neuron (black circle), while attention was directed either to the fixation point (blue dashed circle) or to the contralateral stimulus (red dashed circle). The two stimuli were always moving in the same direction. (Middle) Tuning functions of a motion-selective MT neuron. (Right) Simulations of the normalization model. Reproduced from Reynolds and Heeger (2009).

2 *Attention as noise reduction*

While prior probability can explain some aspects of attention, it's unlikely to be a complete account (see Whiteley and Sahani, 2012). The basic issue is that not all attentional phenomena make sense as changes in the prior. When multiple stimuli are presented simultaneously and the subject has been trained to respond to only one of them, we can speak of attention to the response-relevant stimulus, but that stimulus does not necessarily have higher probability in terms of its location or feature value. What's needed is a different conceptualization of attentional effects that captures allocation of limited/costly cognitive resources.

There is a long history of ideas about capacity limits, resource allocation, and attention. See for example the book by Kahneman (1973).

2.1 Optimizing signal precision

We will analyze a setup in which the agent receives a set of N independent Gaussian signals with mean s and precision (inverse variance) λ, summarized by an aggregate signal $x \sim \mathcal{N}(s, 1/(N\lambda))$. The stimulus is drawn from a prior $s \sim \mathcal{N}(\bar{s}, 1/\lambda_0)$. This is just a version of the Gaussian magnitude estimation problem that we already encountered in Chapter 4. Recall that the posterior mean $\hat{s}$ is a convex combination of the prior mean and the signal:

We will refer to N as the *sample size*, inspired by earlier psychophysical models (Swets et al., 1959; Bonnel and Miller, 1994).

$$\hat{s} = wx + (1 - w)\bar{s}, \tag{12}$$

where

$$w = \frac{N\lambda}{\lambda_0 + N\lambda} \tag{13}$$

is the *signal sensitivity*. The new twist is that we will now allow the agent to adjust the signal precision—the allocation of attention is a kind of cognitive action which the agent can optimize (Gershman and Burke, 2023).

To accommodate a wide variety of tasks with different utility functions, we generalize the setup in Chapter 4, assuming that the utility is a monotonically decreasing and differentiable function of the squared error $\epsilon = (s - \hat{s})^2$. This allows us to approximate it with a Taylor series:

$$u(\epsilon) \approx u(0) - \beta(s - \hat{s})^2, \tag{14}$$

where $u(0)$ is the maximum achievable utility. and

$$\beta = -u''(0) > 0 \tag{15}$$

is the *attentional incentive*, which determines the degree to which utility is contingent on error. Intuitively, agents should be more motivated to pay attention when this contingency is stronger.

The expected utility is given by:

$$\bar{u}(\lambda) = \mathbb{E}[u(\epsilon)|\lambda] \approx u(0) - \frac{\beta}{N\lambda + \lambda_0}. \tag{16}$$

Thus, the agent achieves higher utility when: (i) the attentional incentive is smaller (i.e., the agent doesn't need to pay attention to earn utility); (ii) the signal precision is larger; (iii) the prior precision is larger; and (iv) the sample size is larger.

Next, we need to specify the cost of attention. Recall from Chapter 4 how we used the KL divergence between the (approximate) posterior and the prior as a cost function, which penalizes large belief updates. We noted that when inference is exact, the expected cost corresponds to the mutual information between the hidden state and the signal. For the Gaussian generative model, the expected cost can be expressed in closed form:

$$c(\lambda) = \frac{1}{2} \log \left(1 + \frac{N\lambda}{\lambda_0} \right). \tag{17}$$

We are now ready to state the optimization problem:

$$\lambda^* = \underset{\lambda}{\operatorname{argmax}}\, \bar{u}(\lambda) - \kappa c(\lambda), \tag{18}$$

where $\kappa > 0$ is an *attentional cost* parameter that captures the agent's limited capacity. The solution to this problem is an attentional policy that outputs the optimal signal precision:

In the economics literature, this kind of optimization problem is known as *rational inattention* (Sims, 2003). It was originally developed to explain why agents under-react to market movements, based on the idea that information is cognitively costly.

$$\lambda^* = \max(0, (2\beta/\kappa - \lambda_0)/N). \tag{19}$$

The optimal signal precision grows with the attentional incentive and shrinks with the attentional cost. In addition, it is smaller for higher prior precision, due to the fact that the agent doesn't need to pay attention to the signal as much when they are more confident prior to observing the signal. Finally, all of these factors are attenuated with a larger sample size; the same amount of information can be obtained with less attention by observing for longer.

Plugging the optimal signal precision back into the expected utility yields:

$$\bar{u}(\lambda^*) \approx u(0) - \frac{\kappa}{2}. \tag{20}$$

Unsurprisingly, expected utility is lower when the expected cost is higher. We can also plug the optimal signal precision into the expression for the sensitivity:

$$w^* = 1 - \frac{\lambda_0 \kappa}{2\beta}. \tag{21}$$

The agent will be more responsive to signals when: (i) the prior precision is lower; (ii) the attentional cost is lower; and (iii) the attentional incentive is higher.

2.2 *Detection tasks*

In a standard detection task, a subject is presented with a signal x and needs to judge whether a stimulus was present or absent. We will assume that the "stimulus present" (h_1) distribution corresponds to $p(x|h_1) = \mathcal{N}(x; s, \sigma^2)$, and the "stimulus absent" (h_0) distribution corresponds to $p(x|h_0) = \mathcal{N}(x; 0, \sigma^2)$. In other words, stimulus absence generates a noisy signal centered at 0 with a precision matched to stimulus presence. One can reasonably question these assumptions, but we'll stick with them for now since they facilitate analytical calculations.

To solve the detection task optimally, an agent should choose h_1 when the log-odds $\log \frac{p(x|h_1)}{p(x|h_0)}$ exceeds a decision threshold B. The log-odds is given by:

$$\log \frac{p(x|h_1)}{p(x|h_0)} = \frac{sx}{\sigma^2} - \frac{s^2}{2\sigma^2} \tag{22}$$

Setting this equal to B and solving for x yields a threshold x_B on the signal:

$$x_B = \frac{s}{2} + \frac{\sigma^2 B}{s}. \tag{23}$$

To isolate the role of signal precision in this computation, it is useful to define a *discriminability index*, $d' = s/\sigma$. This quantity can derived from the hit rate H (proportion of stimulus-present trials that were correctly identified) and the false alarm rate FA (proportion of stimulus-absent trials that were misidentified):

$$d' = \Phi^{-1}(\mathrm{H}) - \Phi^{-1}(\mathrm{FA}), \tag{24}$$

where $\Phi^{-1}(\cdot)$ is the inverse Gaussian cumulative distribution function. This equation allows experimenters to empirically estimate the discriminability index and compare it across conditions. Additionally, experimenters can estimate the threshold using the following relation:

$$B = \frac{d'}{2} - \Phi^{-1}(H). \tag{25}$$

The threshold captures any bias a subject might have that is independent of the signal.

These measures of discriminability and decision threshold are important to separate because there is a long history of debate about which processes are affected by attention in detection tasks. A typical manipulation of attention in such tasks is to cue a particular location prior to a signal appearing there. One possibility is that the cue biases responses towards detection (by lowering the decision threshold) without improving discriminability; this would have the effect of increasing both the hit and false alarm rates. Another possibility is that the cue improves discriminability (by increasing signal precision) without changing bias; this would have the effect of increasing the hit rate while reducing the false alarm rate. It's also possible that both or neither changes occur.

Despite decades of experimental work, these questions were not decisively answered due to a host of methodological subtleties (see

The classical body of theory on detection tasks is known as *signal detection theory* (Green and Swets, 1966). This section derives parts of this theory from first principles.

In the usage here, *discriminability* refers to a behaviorally quantifiable measure, whereas *sensitivity* and *precision* refer to cognitive variables that may contribute to changes in discriminability. Note that some treatments of signal detection theory use "sensitivity" to refer to what we are calling here discriminability.

In detection tasks, the decision threshold is also known as the *criterion* or *bias*.

Carrasco, 2006). Some studies have reported changes in discriminability (e.g., Downing, 1988; Carrasco et al., 2000); some have reported changes in the decision threshold (e.g., Müller and Findlay, 1987; Palmer et al., 1993); and some have reported changes in both (e.g., Hawkins et al., 1990). One experimental factor that appears to be critical is whether stimuli are backward-masked—a seemingly innocuous procedural detail, whereby a mask (e.g., a checkerboard pattern) is briefly presented in the stimulus location immediately after stimulus offset. Backward masks are thought to interrupt feedforward visual processing, allowing experimenters to precisely control the duration of early sensory representations. Reviewing the detection literature, Smith and Ratcliff (2009) concluded that most studies reporting a change in discriminability used backward masks. A direct comparison of the same detection task with and without backward masking was undertaken by Smith et al. (2004), providing a clear demonstration that changes in discriminability only appeared with backward masks (Figure 7).

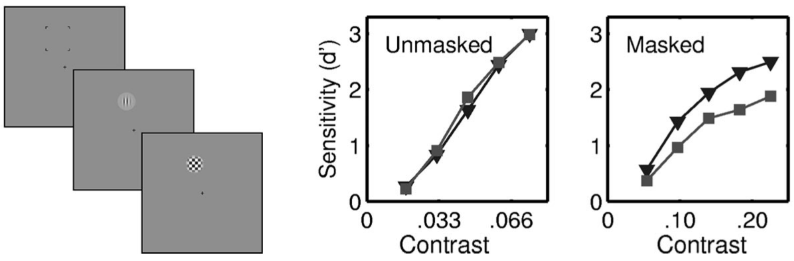

Figure 7: **Changes in discriminability depend on backward masking**. (Left) Experimental paradigm. A stimulus (vertically or horizontally oriented Gabor patch) is probabilistically presented at a cued or uncued location, followed by a backward mask. Subjects were instructed to judge stimulus orientation (a proxy for detection), though other tasks yield similar results. (Right) The discriminability as a function of stimulus contrast for correctly cued stimuli (triangles) and incorrectly cued stimuli (squares). Adapted from Smith et al. (2004).

Smith and Ratcliff suggested that the limited flow of information into visual memory due to backward masking plays a pivotal role in attentional effects. They developed a rather complex formal model of how this might work (based on evidence accumulation concepts reviewed in the next chapter). Here we will use the optimal signal precision model to arrive at a similar conclusion.

One way to model backward masking is by reducing the sample size parameter N, which constrains the amount of sensory information available to further processing. As noted earlier, this has the effect of attenuating all the other factors governing optimal signal precision. Consequently, changes in attentional incentive (e.g., by offering a larger reward for correct responses) or attentional cost (e.g., by directing limited attentional resources to a particular spatial location) have weaker effects on performance, consistent with the empirical literature and the work of Smith et al. (2004).

2.3 *Neural correlates of discriminability and decision threshold modulation*

As pointed out by Luo and Maunsell (2015), many neurophysiological studies of attention conflate effects of attention on discriminability and decision threshold. Consequently, it is challenging to interpret the neural activity in terms of the underlying computational parameters. To remedy this, Luo and Maunsell developed a new variant of the detection task for monkeys which disentangled these components. In some sessions, they lowered the decision threshold at particular locations by increasing rewards for hits relative to correct rejections (responding "absent" when the stimulus is absent). This had the effect of increasing both the hit rate and the false alarm rate at those locations. In other sessions, they lowered discriminability at particular locations by increasing the reward for both hits and correct rejections, which had the effect of increasing the hit rate and decreasing the false alarm rate.

We've already seen how V4 is a particularly common site of attentional modulation in visual tasks. Accordingly, Luo and Maunsell recorded from neurons in this area during the two sessions types. Their key finding was that firing rates were modulated by changes in discriminability but not by changes in the decision threshold. They also found that neural variability in V4 was decreased when discriminability increased. This is consistent with an increase in signal precision.

Attentional effects on the decision threshold may occur in the superior colliculus. Recall from Chapter 4 our discussion of evidence that this region encodes a probabilistic map of visual target locations, which is used to guide eye movements. Shifts in activity of this region may bias eye movements towards particular locations. Consistent with this idea, studies that stimulate or inactivate the superior colliculus induce changes in the decision threshold without affecting discriminability (Sridharan et al., 2017). We will have more to say about the contribution of superior colliculus to perceptual decisions in the next chapter.

2.4 *Dopamine and signal precision*

Several lines of evidence suggest that signal precision is controlled by the neuromodulator dopamine (Friston et al., 2012; Mikhael et al., 2021). Pharmacologically elevating dopamine generally increases sensitivity to stimuli. For example, elevating dopamine using methylphenidate increases sensitivity in a motion discrimination task (Beste et al., 2018). Patients with Parkinson's disease (which dramatically depletes dopamine) exhibit attenuated sensitivity to mechanical stimulation that was ameliorated in proportion to their dopamine medication dose (Wolpe et al., 2018). These findings are consistent with the finding that injection of dopamine into prefrontal areas increases the magnitude and selectivity of responses in visual cortex (Noudoost and Moore, 2011). Dopamine's effects on signal precision may thus be mediated by prefrontal areas that provide the directive signals for attentional enhancement in sensory areas.

This hypothesized role primarily relates to slow timescale (tonic) fluctuations in dopamine levels, distinct from the role of faster timescale (phasic) fluctuations which we will discuss further in Chapter 10. A prominent theory posits that tonic dopamine encodes the average reward rate in the environment (Niv et al., 2007). This theory was motivated by the observation that pharmacological manipulations of dopamine affect response vigor: higher reward rates (and hence higher dopamine) should enhance willingness to work for reward. We can connect the signal precision account to this idea by observing that optimal signal precision will in general be higher when the average reward is higher.

3 *Conclusion*

Attention is not a unitary construct—it consists of several components that affect behavior and neural activity in different ways. Nonetheless, we can explain many aspects of attention modulation within a unified framework. The key concepts are shifts in prior probabilities (e.g., due to spatial cues), shifts in discriminability (e.g., due to optimization of signal precision), and shifts in decision threshold (e.g., by rewarding hits more than correct rejections). These attentional shifts are linked to distinct neural correlates. However, it remains unclear how these correlates interact to produce the relevant psychophysical changes.

Study questions

1. Why does attention sometimes multiplicatively enhance tuning curves and sometimes sharpen them?
2. Which empirical phenomena are better explained in terms of prior probability vs. noise reduction?
3. Compare and contrast the computational roles of acetylcholine and dopamine in attention.

References

Beste, C., Adelhöfer, N., Gohil, K., Passow, S., Roessner, V., and Li, S.-C. (2018). Dopamine modulates the efficiency of sensory evidence accumulation during perceptual decision making. *International Journal of Neuropsychopharmacology*, 21:649–655.

Bonnel, A.-M. and Miller, J. (1994). Attentional effects on concurrent psychophysical discriminations: Investigations of a sample-size model. *Perception & Psychophysics*, 55:162–179.

Bowman, E. M., Brown, V. J., Kertzman, C., Schwarz, U., and Robinson, D. L. (1993). Covert orienting of attention in macaques. I. Effects of behavioral context. *Journal of Neurophysiology*, 70:431–443.

Carrasco, M. (2006). Covert attention increases contrast sensitivity: Psychophysical, neurophysiological and neuroimaging studies. *Progress in Brain Research*, 154:33–70.

Carrasco, M., Penpeci-Talgar, C., and Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. *Vision Research*, 40:1203–1215.

Chikkerur, S., Serre, T., Tan, C., and Poggio, T. (2010). What and where: A Bayesian inference theory of attention. *Vision Research*, 50:2233–2247.

Chun, M. M., Golomb, J. D., and Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, 62:73–101.

Davidson, M., Cutrell, E., and Marrocco, R. (1999). Scopolamine slows the orienting of attention in primates to cued visual targets. *Psychopharmacology*, 142:1–8.

Davidson, M. C. and Marrocco, R. T. (2000). Local infusion of scopolamine into intraparietal cortex slows covert orienting in rhesus monkeys. *Journal of Neurophysiology*, 83:1536–1549.

Dayan, P. and Zemel, R. S. (1999). Statistical models and sensory attention. In *1999 Ninth International Conference on Artificial Neural Networks ICANN 99.(Conf. Publ. No. 470)*, volume 2, pages 1017–1022. IET.

Downing, C. (1988). Expectancy and visual-spatial attention: effects on perceptual quality. *Journal of Experimental Psychology. Human Perception and Performance*, 14:188–202.

Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E., and Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Computational Biology*, 8:e1002327.

Gershman, S. J. and Burke, T. (2023). Mental control of uncertainty. *Cognitive, Affective, & Behavioral Neuroscience*, 23:465–475.

Green, D. M. and Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*. Wiley New York.

Hasselmo, M. E. and Sarter, M. (2011). Modes and models of forebrain cholinergic neuromodulation of cognition. *Neuropsychopharmacology*, 36:52–73.

Hawkins, H., Hillyard, S., Luck, S., Mouloua, M., Downing, C., and Woodward, D. (1990). Visual attention modulates signal detectability. *Journal of experimental psychology. Human perception and performance*, 16:802–811.

Kahneman, D. (1973). *Attention and Effort*. Prentice-Hall.

Luo, T. Z. and Maunsell, J. H. (2015). Neuronal modulations in visual cortex are associated with only one of multiple components of attention. *Neuron*, 86:1182–1188.

Martinez-Trujillo, J. C. and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current biology*, 14:744–751.

McAdams, C. J. and Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, 19:431–441.

Mikhael, J. G., Lai, L., and Gershman, S. J. (2021). Rational inattention and tonic dopamine. *PLoS Computational Biology*, 17:e1008659.

Müller, H. J. and Findlay, J. M. (1987). Sensitivity and criterion effects in the spatial cuing of visual attention. *Perception & Psychophysics*, 42:383–399.

Niv, Y., Daw, N. D., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*, 191:507–520.

Noudoost, B. and Moore, T. (2011). Control of visual cortical signals by prefrontal dopamine. *Nature*, 474:372–375.

Palmer, J., Ames, C., and Lindsey, D. (1993). Measuring the effect of attention on simple visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 19:108–130.

Phillips, J. M., McAlonan, K., Robb, W. G., and Brown, V. J. (2000). Cholinergic neurotransmission influences covert orientation of visuospatial attention in the rat. *Psychopharmacology*, 150:112–116.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32:3–25.

Posner, M. I., Walker, J. A., Friedrich, F. J., and Rafal, R. D. (1984). Effects of parietal injury on covert orienting of attention. *Journal of Neuroscience*, 4:1863–1874.

Rao, R. P. (2005). Bayesian inference and attentional modulation in the visual cortex. *Neuroreport*, 16:1843–1848.

Reynolds, J. H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19:1736–1753.

Reynolds, J. H. and Heeger, D. J. (2009). The normalization model of attention. *Neuron*, 61:168–185.

Reynolds, J. H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, 26:703–714.

Robinson, D. L., Bowman, E. M., and Kertzman, C. (1995). Covert orienting of attention in macaques. II. Contributions of parietal cortex. *Journal of Neurophysiology*, 74:698–712.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50:665–690.

Smith, P. L. and Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review*, 116:283–317.

Smith, P. L., Ratcliff, R., and Wolfgang, B. J. (2004). Attention orienting and the time course of perceptual decisions: Response time distributions with masked and unmasked displays. *Vision Research*, 44:1297–1320.

Spitzer, H., Desimone, R., and Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science*, 240:338–340.

Sridharan, D., Steinmetz, N. A., Moore, T., and Knudsen, E. I. (2017). Does the superior colliculus control perceptual sensitivity or choice bias during attention? evidence from a multialternative decision framework. *Journal of Neuroscience*, 37:480–511.

Stewart, C., Burke, S., and Marrocco, R. (2001). Cholinergic modulation of covert attention in the rat. *Psychopharmacology*, 155:210–218.

Swets, J. A., Shipley, E. F., McKey, M. J., and Green, D. M. (1959). Multiple observations of signals in noise. *The Journal of the Acoustical Society of America*, 31:514–521.

Whiteley, L. and Sahani, M. (2012). Attention in a Bayesian framework. *Frontiers in Human Neuroscience*, 6:100.

Williford, T. and Maunsell, J. H. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, 96:40–54.

Wolpe, N., Zhang, J., Nombela, C., Ingram, J. N., Wolpert, D. M., and Rowe, J. B. (2018). Sensory attenuation in Parkinson's disease is related to disease severity and dopamine dose. *Scientific Reports*, 8:15643.

Yu, A. and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46:681–692.