

Learning and selective attention

Peter Dayan¹, Sham Kakade¹ and P. Read Montague²

¹ Gatsby Computational Neuroscience Unit, University College London, 17 Queen Square, London WC1N 3AR, UK

² Division of Neuroscience, Baylor College of Medicine, 1 Baylor Plaza, Houston, Texas 77030, USA

Correspondence should be directed to P.D. (dayan@gatsby.ucl.ac.uk)

Selective attention involves the differential processing of different stimuli, and has widespread psychological and neural consequences. Although computational modeling should offer a powerful way of linking observable phenomena at different levels, most work has focused on the relatively narrow issue of constraints on processing resources. By contrast, we consider statistical and informational aspects of selective attention, divorced from resource constraints, which are evident in animal conditioning experiments involving uncertain predictions and unreliable stimuli. Neuromodulatory systems and limbic structures are known to underlie attentional effects in such tasks.

Attention is a complex neural and psychological phenomenon, coming in many different forms and involving many different brain structures and mechanisms. Oft-quoted examples include arousal, vigilance and selection. Most empirical studies of attention concentrate on the nature, control and consequences of selection. From the outset^{1,2}, models of selective attention have largely focused on the idea that there might be limited computational resources available to process inputs and choose and execute courses of action. Selective attention is usually conceived as a response to these constraints, in which all but the most important or relevant stimuli for a task are filtered out³. This might occur in conjunction with a serial processing strategy in which different stimuli are selectively attended in turn, and the results of processing each are stored in working memory.

Computational models might be expected to link together the rich set of psychological and neural findings on selective attention, while providing an overall rationale for what is observed. However, the focus on the single issue of resource constraints leaves open at least two concerns. First, although animals must have a computational strategy to deal with their limited resources, is this the only rationale for selective attention? One reason to doubt this is that attentional effects are evident in tasks without abundant complex stimuli, in which input filtering would seem unnecessary. Second, few constraints on how selection should work come merely because it is necessary to overcome resource limitations. Altogether, we need a computationally richer conception of selection that does not depend only on resource constraints, and which helps formalize what it means to be an important and relevant stimulus.

Here we show some work in this direction, by considering cases in which selective attention is computationally sensible by itself, and is not a response to any resource constraint⁴. Although the general approach has also been applied to visual attentional tasks⁵, we illustrate it here through examples of attention in classical conditioning⁶. Classical conditioning is convenient because only a few stimuli are usually present at any one time, and their significance (and therefore how much they might be selected) can be manipulated in a graded manner. Further, the neural substrate underlying performance in attentional tasks in conditioning has been probed using selective lesions. Computational models of attention in conditioning are primarily statistical models, concerned with identifying

and assessing the relevance of stimuli for predicting outcomes, and the relevance of failure to predict outcomes correctly for adjusting the predictions.

Classical conditioning

Classical conditioning is the study of how animals learn predictive relationships⁷⁻⁹. In their natural environments, animals face a multitude of stimuli, very few of which are likely to be useful as predictors. In keeping with ideas of resource constraints, various theories of classical conditioning have appealed to a limited-capacity learning processor. These theories specify rules governing which stimuli are deemed predictively useful and thus gain access to this processor by being selectively attended^{10,11}. Instead, we consider quantitatively how statistical models of learning demand selection as a result of optimal processing. Thus we can abandon the notion of capacity limitations, as they are redundant⁴. We consider and provide simple statistical models for two different sorts of attentional selection that emerge in conditioning¹²⁻¹⁴. One concerns how learning should be competitively allocated among stimuli^{10,11}; the other concerns how responsibility for making predictions should be competitively allocated⁹.

Behavioral studies show that animals do indeed pay different amounts of attention to different stimuli. In rats, experimental evidence implicates neuromodulatory systems in such selective attention¹⁵⁻¹⁷. In appetitive conditioning, cholinergic projections from the basal forebrain to the parietal cortex and hippocampal formation are important in what we identify as selection for learning, in which certain stimuli change their predictions faster than other stimuli. We suggest that the nucleus accumbens is similarly important in the selection of which stimuli to use in prediction.

In a standard protocol for appetitive classical conditioning (Fig. 1), a hungry animal is faced with numerous stimuli. Some stimuli are directly controlled by the experimenter (blue, red and green lights), and one stimulus (yellow) is uncontrolled, for instance representing noise sources from the room, random fluctuations in lighting and so on. As best the animal knows when the experiment begins, all these 'conditioned stimuli' might bear some relationship to the delivery of food pellets, the 'unconditioned stimulus'. A common way⁷ of interpreting classical conditioning is that animals learn predictive rela-

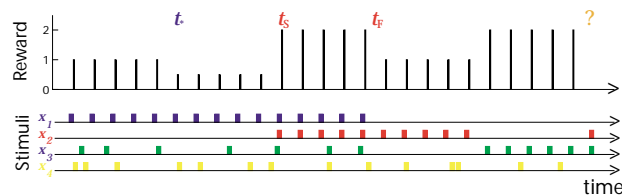


Fig. 1. Uncertainty and unreliability. Top, rewards of various magnitudes (size or number of food pellets) are given. Bottom, lights of different colors represented by binary $x_i(t)$ potentially predict reward. Time t_* , the first time the reward associated with the blue stimulus changes; time t_s , when the red light is introduced; time t_f , the end of the plots in Fig. 2; time '?', when the prediction associated with the combination of red and green lights is assessed.

tionships in their environments, that is, they learn what the conditioned stimuli predict about the unconditioned stimulus. Prediction of important outcomes, such as food or aversive unconditioned stimuli (for example, electric shocks), is a very important task likely to underlie the animals' capacity to choose behaviors appropriate to their environments. Their predictions can be probed in various ways, including monitoring their behavior for orienting responses to conditioned stimuli that predict food, approach movements toward the places where they are fed, or, when the conditioned stimuli predict shocks, suppression of a specific active behavior (such as lever pressing for food) while an aversive conditioned stimulus is presented.

Such experiments, and also more naturalistic prediction tasks, have several important features in common. First, only some conditioned stimuli predict the outcome consistently or reliably throughout the experiment (Fig. 1, blue and red lights). Second, the experimenter changes the contingencies during the experiment, and so the predictions must change too. To deal with this, animals need to use the two types of selective attention mentioned above. Note in Fig. 1 that the amount of reward that follows the blue light changes twice. The first change is not contingent on the presentation of any new stimulus (at time t_*), and the second change coincides with the onset of a novel red light (at time t_s).

Any system, natural or artificial, that estimates the current predictive relationship between conditioned stimuli and reward based on past observations can be described as making statistical inferences. The rectitude of such inferences depends on the system's (possibly implicit) underlying statistical model of the world. In common with other rational analyses, we make the critical assumption that the animals make correct, or at least approximately correct, inferences. Then, we can use a two-pronged approach to understand the results of conditioning experiments, regarding performance in conditioning experiments as probing the actual statistical models underlying animal inferences, and considering statistical models that take sensible note of unreliable and changing predictive relationships. This approach leads to the two models of selective attention we discuss below.

Uncertainty

Box 1 describes a simple statistical model of the animal's task in classical conditioning, in terms of a contrivance called a 'Kalman filter'^{18,19}. This model formalizes the predictive relationship between conditioned stimuli and reward, together with how this predictive relationship is expected to change over time. The model thereby also formalizes how the animal should make inferences about the present predictive relationship from past observations. Importantly, attached to the prediction made for each stimulus is the degree of uncertainty in the prediction. The uncertainty is large if the stimulus has not been shown in the recent past. The uncertainty is small if the animal has recently seen substantial evidence about the relationship between that stimulus and the outcome.

In the Kalman filter, predictions of a stimulus are changed when the outcome is predicted incorrectly, based on a compromise between prior information and recent observations. For instance, more prior observations of the relationship between a stimulus and reward generally mean that this relationship is more certain. Therefore, the animal should be slower to change its estimate of this relationship based on a prediction error. However, if the error in prediction coincides with a novel stimulus, whose relationship with the reward is uncertain, then this relationship should change quickly.

For example, at time t_* (Fig. 1), the animal might expect one pellet of food, based on its recent experience with the blue light. However, it actually receives half a pellet. The resulting prediction error drives learning. Under the Kalman filter model (Fig. 2), up to time t_f in Fig. 1, the blue light increases its prediction to about one pellet up to time t_* , then decreases its prediction because of the prediction error (Fig. 2a). The speed at which the prediction of the blue light adapts to half a pellet depends on the uncertainty of the blue light's predictions (Fig. 2c).

Responsibility for the prediction error is competitively allocated between the stimuli present in inverse proportion to the uncertainties associated with each of their predictions (Box 1). This competitive allocation governs how much the predictions of each stimulus are changed. For instance, at time t_s in Fig. 1, the animal might expect half a pellet of food, based on its recent experience with the blue light and lack of prior experience with the red light. However, in response to the predictive failure at this time, the prediction associated with the red light will change substantially more than the prediction associated with the blue light. This happens because the animal is much less certain about the association between the red light and

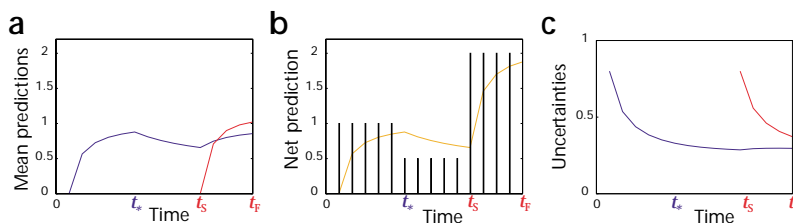


Fig. 2. Prediction with the Kalman filter up to time t_f . (a) Mean value of predictions associated with the blue and red lights. (b) The net prediction (brown) of the actual reward (black) that would be made according to the sum of predictions of the stimuli. Note how the prediction adapts to the change in the actual rewards. (c) Uncertainties associated with the blue and red lights. Larger uncertainties (as at the start of learning for each stimulus) allow larger changes in the predictions. Note that there is asymptotic uncertainty, that is, the estimates do not become arbitrarily accurate (because of the continual possibility of change in the world). Predictions and uncertainties for the green and yellow lights are not shown.

Box 1. Conditioning and the Kalman filter

We first consider a simple statistical model for conditioning¹⁸. The presence and absence of the conditioned stimuli (the lights) is represented by a vector $\mathbf{x}(t)$, with one component for each stimulus. Here, $x_i(t) = 1$ if stimulus i is present at time or trial t , and is 0 otherwise. Similarly, the delivery of the unconditioned stimulus (the food pellets) is represented by the scalar $r(t)$. In the statistical model, there is a true relationship in the world between the lights and the food represented by a set of parameters $\mathbf{w}(t)$. These correspond to the way that the experimenter has programmed the apparatus. These parameters can change over time, for instance in the transitions between the different regimes apparent in Fig. 1. In the simplest case, we might write

$$r(t) = \mathbf{x}(t) \cdot \mathbf{w}(t) + \varepsilon(t)$$

where $\mathbf{x}(t) \cdot \mathbf{w}(t) = \sum_j x_j(t) w_j(t)$ indicates that the weights associated with the stimuli are combined additively (a point to which we return when we consider reliability competition). The term $\varepsilon(t)$ represents noise, either in the actual delivery of the reward or in the animal's processing of its magnitude.

The most straightforward model of how $\mathbf{w}(t)$ might change over time is

$$\mathbf{w}(t+1) = \mathbf{w}(t) + \eta(t)$$

where $\eta(t)$ is a drift term, the details of which are unknown to the animal. The inference task for the animal is to take the paired sequences of observations $\{\mathbf{x}(t), r(t)\}$ and infer $\mathbf{w}(t)$. That is (Fig. 3), given sequences such as those shown in Fig. 1, provide a running estimate, $\hat{\mathbf{w}}(t)$, of the relationship between

reward (about which nothing has previously been observed) than that between the blue light and reward. At time t_5 , the prediction associated with the red light changes more than that associated with the blue light (Fig. 2a). The net prediction made on the basis of both stimuli (Fig. 2b) approaches two pellets (the net prediction being, in this case, the sum of the predictions of the individual stimuli). The uncertainties in the two predictions (Fig. 2c), which decrease with more observations, determine this difference. The asymptotic level of this uncertainty governs the rate at which the predictions can change once fully established, which in turn is determined by the stability of the environment, that is, how fast the actual reward contingencies should be expected to change.

The competitive allocation of learning between the two lights is exactly a form of selective attention for learning. It arises not from a resource limitation, but rather from a normative treatment of the statistical contingencies involved in prediction, in which the predictions are optimally inferred in a changing world as specified by the Kalman filter. This form of attention is closely related to the Pearce-Hall attentional theory of conditioning¹¹, from which the Kalman filter model borrows the idea that learning should depend on predictive uncertainty, so that more uncertain predictions change more quickly. Under the Pearce-Hall theory, for which there is substantial empirical evidence^{20–23}, the degree of uncertainty about each stimulus is determined by its past failures in prediction (for instance, at time t_* for the blue stimulus) rather than just past observations, as also emerges in more sophisticated versions of the Kalman filter model¹⁸.

stimuli $\mathbf{x}(t)$ and reward $r(t)$, and thus a running estimate or prediction of the reward $\hat{r}(t) = \mathbf{x}(t) \cdot \hat{\mathbf{w}}(t)$. If just one stimulus is presented (say, the blue light), then the prediction formed by the animal will reveal the association learned between the blue light and reward. In the simplest case, the running estimate of the mean prediction associated with stimulus i , $\hat{w}_i(t)$, is accompanied by a running estimate of the uncertainty in that prediction, $\sigma_i^2(t)$. This reflects a balance between the information acquired from the observations, and the drift term of Eq. 2, which permits predictive relationships to change over time.

Making simplifying assumptions, it can be shown¹⁸ that adaptation of $\hat{\mathbf{w}}(t)$ should be based on the formulae

$$\begin{aligned} \hat{w}_i(t+1) &= \hat{w}_i(t) + \alpha_i(t) \delta(t) & \delta(t) &= r(t) - \mathbf{x}(t) \cdot \hat{\mathbf{w}}(t) \\ \alpha_i(t) &= \frac{\sigma_i^2(t) x_i(t)}{\sum_j \sigma_j^2(t) x_j(t) + E} \end{aligned}$$

where E is the variance of the noise $\varepsilon(t)$ corrupting $r(t)$. Here, the term $\delta(t)$ reports the error in the current predictions⁴⁹ and is exactly the same error term that appears in many supervised learning rules, from the Rescorla-Wagner⁵⁰ rule for conditioning, to the backpropagation learning rule for neural networks, and, in slightly modified form, the temporal difference error⁵¹ associated with the activity of midbrain dopamine cells⁵². The associability term $\alpha_i(t)$, which governs the learning rate, contains the critical feature of a competitive allocation of learning between the stimuli according to their uncertainties—it is as if the prediction error is divided out among all the stimuli that are present according to their individual relative uncertainties. The Kalman filter model also specifies how $\sigma_i^2(t)$ changes over time.

Unreliability

Under the model above, the predictions made for all stimuli at a particular time are simply added. There is evidence that such an integration mechanism is incomplete and in some instances incorrect²⁴. For example, consider the time marked by '?' in Fig. 1. In the recent past, the red light has been associated with delivery of one pellet of food, and the green light with two pellets of food. Further, the two stimuli have not recently been presented together. If these two lights are now both presented, what

Fig. 3. Architecture for prediction. The stimuli $\mathbf{x}(t)$ are represented by the activity of the input nodes and act through the weights $\hat{\mathbf{w}}(t)$ (and, later, a competitive combination rule) to construct a prediction $\hat{r}(t)$ of the reward $r(t)$. Weight adaptation is based on the prediction error $\delta(t)$. We model the weights as being in the basolateral nuclei of the amygdala, competitive combination in the nucleus accumbens, the report of prediction error for appetitive unconditioned stimuli coming from the dopamine system⁴⁹, and control over plasticity (not shown) exerted indirectly by the central nucleus of the amygdala and hippocampus.





Fig. 4. Competitive combination of predictions. Red and green lines, example distributions of predictions of reward from the red and green lights at the time marked by '?' in Fig. 1. The prediction of the red light is more reliable, as its distribution is more sharply peaked. Brown line, combined prediction after information from both stimuli has been integrated. As the red light is more reliable, the mean of the combined estimate is closer to w_2 .



should the net prediction of reward be? Arguably, it should be one, two or three pellets, or indeed something in between, depending on how competitively or cooperatively the predictions are combined.

Competitive combination^{9,24} arises when the final prediction of the reward averages the predictions associated with the conditioned stimuli that are present in such a way that the predictions of some stimuli contribute more than those of others. In our example, an obvious basis for competition is that, for most of the experiment, the green light has been explicitly unrelated to reward. Therefore, the animal might consider it unreliable as a predictor. Conversely, the red light has been a consistently good predictor of reward. Given its unreliability, if only the green light were presented, then, short of any other information, the best prediction might be two pellets. However, if both the red and green lights were presented, the best combined prediction is more likely closer to one pellet, the value associated with the red light alone.

In a statistical model of this (**Box 2**), stimuli compete to participate in a combined prediction in proportion to their estimated reliability as predictors. These reliabilities are expected^{14,25} to be adapted over a longer time scale than the predictions themselves. Competition between stimuli based on their reliabilities is another form of selective attention. It is also not induced by any resource constraint, but is a normative consequence of a natural statistical model for prediction. It is possible (though messy) to integrate the models of uncertainty and unreliability. In particular, reliabilities (which lead to competition for making the predictions) are different from uncertainties in prediction (which lead to competition for learning the predictions). Reliabilities are the statistical account of which stimuli the animal deems important in prediction, whereas uncertainties quantify how well those predictions are known. One can have a reliable predictor that is clearly important, like the red light, but whose current prediction is uncertain because it is based on little information.

DISCUSSION

By considering simple forms of attention in classical conditioning, we have shown that selection can be an appropriate solution to an inference problem, rather than just a way of addressing a resource constraint. Selection is useful if the demands of a task imply that some available information is more relevant than other information (just as some available stimuli are more reliable predictors than others). This sort of selection is rarely absolute: only in special circumstances should one pay no regard whatsoever to a stimulus; rather, some stimuli will be relatively downgraded with respect to others. We presented simple statistical models in which such selection is normative; other, more complex statistical models would lead to similar conclusions²⁶.

We can compare this statistical conception of selective attention with two associated views. One, a computational idea inspired by neurophysiological data on the effect of attention on receptive fields, is that selection should eliminate all traces of the unattended stimuli, leaving just the activity associated with the attended stimulus as if it had been presented alone²⁷. Models show that this can indeed result from the forms of attention we have considered. Of course, input information about the stimulus on which the neural activity is based may be different in the face of attention^{5,28} (for instance, based on only the relevant part of a cell's receptive field). In this case, one might expect to find that the activity will also be different, for instance stronger^{29,30}.

The second associated view, which is more mechanistic, is that neuromodulators might change the signal-to-noise ratio of single cells³¹ by suppressing apparently spontaneous activity, compared with stimulus-driven activity. This view seems often misunderstood. If cells can operate at a high signal-to-noise ratio, then it would seem strange that they do not always do so. It would seem especially strange if the high signal-to-noise ratio can be achieved by decreasing apparently spontaneous activity without increasing the net firing rate. Under a statistical account

Box 2. Unreliability

To formalize the notion that some stimuli are more reliable than others, we need to postulate a model^{53–55} that allows a variable relationship between the true value of the reward $r(t)$ and the prediction $w_i(t)$ associated with each stimulus i . In one such model, each prediction is expected to vary according to a Gaussian distribution with standard deviation $\tau_i(t)$ about $r(t)$. The prediction $w_i(t)$ should be closer to $r(t)$ for a stimulus with a small value of $\tau_i(t)$ than one with a big value, and we therefore define $\rho_i(t) = 1/\tau_i^2(t)$ as the reliability of stimulus i . The probability distributions associated with the predictions of red and green lights are shown in Fig. 4. The red light is more reliable, as is evident from its sharper distribution.

Again, under some simplifying assumptions, the net prediction based on all the stimuli present is

$$\sum_i \pi_i(t) w_i(t)$$

where $\pi_i(t)$ is the combination weight:

$$\pi_i(t) = \frac{\rho_i(t) x_i(t)}{\sum_j \rho_j(t) x_j(t)}$$

which is a weighted average based on the reliabilities $\rho_i(t)$. The combination weight term represents a competitive allocation of predictions among stimuli according to their reliabilities—stimuli that are more reliable have more say in the prediction. Again, these reliability terms define a form of selective attention through this competitive combination. This form of attention can also be combined with the uncertainty attention discussed above¹⁴.

The brown line in Fig. 4 shows the net prediction at the time marked '?' in Fig. 1, after the predictions of all the stimuli have been considered. The net prediction is calculated as the product of the two distributions, normalized so that the probabilities sum to 1. Note that the center of the brown distribution is closer to that of the red light, as expected, and is sharper than those of both the red and green lights, as both contribute information.

like the one we have developed here, the higher signal-to-noise ratio really comes from a computational strategy of selecting from the inputs just those that are appropriate for a particular task (eliminating, for instance, the effects of uncontrolled stimuli like the yellow input in Fig. 1). Different tasks lead to different assessments about which inputs are relevant.

Lesion studies on rats have probed the neural substrates underlying increases and decreases in competitive allocation of learning to stimuli that become newly predictive or unpredictable^{15–17,32}. The level of allocation is often called the associability of the stimulus. The increase in associability associated with failure in prediction that would emerge from a more complex version of the Kalman filter model involves a pathway from the central nucleus of the amygdala to cholinergic cells in the basal forebrain that project to the parietal cortex, itself implicated in spatial attention^{33–35}. Decreases in associability that occur as a stimulus is found to have either no consequence ('latent inhibition'³⁶) or a consistent consequence, are affected by manipulations of the hippocampus and related structures^{16,37,38}, and are also suggested to involve neuromodulatory systems³⁹. Of course, stimulus uncertainties are likely to be subject to multiple influences. For instance, one factor underlying the uncertainties should be novelty (which may actually be assessed by hippocampal structures⁴⁰); stimuli that are novel in a context should automatically be able to attract learning because their predictive relationships with the environment are clearly unknown.

There is less information about the neural substrate underlying the reliability-based selective attention that controls competitive combination of predictions from multiple stimuli. One attractive possibility is that the ventral striatum and its associated basal ganglia structures are involved in attentional competition in a way that parallels the putative involvement of the dorsal striatum and its associated basal ganglia structures in competition between actions^{41–44}. Stimulus reliabilities would then be stored in corticostriatal connections, whereas affective values (predictions associated with stimuli) would be stored in the basolateral nuclei of the amygdala and perhaps orbitofrontal cortex^{45–47}.

Classical conditioning is an attractive test-bed for theories of attention because the requirement for selection (statistical uncertainty or unreliability) itself implies what its effects should be (better learning to predict error or better participation in making predictions). In other circumstances, the same ideas apply, even if the requirements and effects may be more distantly related. For instance, we have suggested that it is helpful to think about the informational constraint of selection, that is, to ask when it is wise to throw away what might look like available information. Under a computational characterization of the visual system in which it is supposed to build a veridical representation of all the visually signalled features of scenes, the answer is essentially never. Selective attention would always seem to be a regrettable necessity. However, under characterizations like that of active or 'animate' vision⁴⁸, in which the goal is to get answers to particular questions posed about the visual world, such selection of aspects of the input appropriate to the question at hand is inescapable. This opens up our whole scheme of inquiry.

ACKNOWLEDGEMENTS

We are grateful to Nathaniel Daw, Anthony Dickinson, Alex Kacelnik, Theresa Long, Terry Sejnowski, and Rich Sutton for discussions covering many of these issues, and particularly to Nathaniel Daw for comments on an earlier version of this paper. The authors' work is funded by the Gatsby

Charitable Foundation (P.D. and S.K.), the National Science Foundation (S.K.) and NIDA (P.R.M.).

RECEIVED 12 JUNE; ACCEPTED 2 OCTOBER 2000

- Broadbent, D. *Perception and Communication* (Pergamon, Elmsford, New York, 1958).
- Kahneman, D. *Attention and Effort* (Prentice-Hall, Englewood Cliffs, New Jersey, 1973).
- Bundesen, C. in *Converging Operations in the Study of Visual Selective Attention* (eds. Kramer, A. F., Coles, M. G. H. & Logan, G. D.) 1–44 (Am. Psychol. Assoc., Washington, DC, 1996).
- Allport, A. in *Attention and Performance* vol. 14 (eds. Meyer, D. E. & Kornblum, S.) 183–218 (MIT Press, Cambridge, Massachusetts, 1993).
- Dayan, P. & Zemel, R. S. Statistical models and sensory attention. *Intl. Conf. Artificial Neural Networks* (1999).
- Robbins, T. W. in *The Attentive Brain* (ed. Parasuraman, R.) 189–220 (MIT Press, Cambridge, Massachusetts, 1998).
- Dickinson, A. *Contemporary Animal Learning Theory* (Cambridge Univ. Press, Cambridge, 1980).
- Mackintosh, N. J. *Conditioning and Associative Learning* (Oxford Univ. Press, Oxford, 1983).
- Grossberg, S. Processing of expected and unexpected events during conditioning and attention: A psychophysiological theory. *Psychol. Rev.* **89**, 529–572 (1982).
- Mackintosh, N. J. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychol. Rev.* **82**, 276–298 (1975).
- Pearce, J. M. & Hall, G. A model for Pavlovian learning: Variation in the effectiveness of conditioned but not unconditioned stimuli. *Psychol. Rev.* **87**, 532–552 (1980).
- Kruschke, J. K. Relating Mackintosh's (1975) theory to connectionist models and human categorization. Talk presented at the *Eighth Australasian Mathematical Psychology Conference* (Perth, Australia, 1997).
- Dayan, P. & Long, T. in *Advances in Neural Information Processing Systems* vol. 10 (eds. Jordan, M. I., Kearns, M. A. & Solla, S. A.) 117–123 (MIT Press, Cambridge, Massachusetts, 1998).
- Kakade, S. & Dayan, P. in *Advances in Neural Information Processing Systems* vol. 12 (eds. Solla, S. A., Leen, T. K. & Muller, K.-R.) 24–30 (MIT Press, Cambridge, Massachusetts, 2000).
- Everitt, B. J. & Robbins, T. W. Central cholinergic systems and cognition. *Annu. Rev. Psychol.* **48**, 649–684 (1997).
- Holland, P. C. Brain mechanisms for changes in processing of conditioned stimuli in Pavlovian conditioning: Implications for behavior theory. *Anim. Learn. Behav.* **25**, 373–399 (1997).
- Holland, P. C. & Gallagher, M. Amygdala circuitry in attentional and representational processes. *Trends Cogn. Sci.* **3**, 65–73 (1999).
- Sutton, R. in *Proceedings of the 7th Yale Workshop on Adaptive and Learning Systems* 161–166 (Yale University, New Haven, Connecticut, 1992).
- Anderson, B. D. O. & Moore, J. B. *Optimal Filtering* (Prentice-Hall, Englewood Cliffs, New Jersey, 1979).
- Pearce, J. M., Wilson, P. N. & Kaye, H. The influence of predictive accuracy on serial conditioning in the rat. *Q. J. Exp. Psychol. Comp. Physiol. Psychol.* **40**, 181–198 (1988).
- Swan, J. A. & Pearce, J. M. The influence of predictive accuracy on serial auto-shaping: Evidence of orienting responses. *J. Exp. Psychol. Anim. Behav. Processes* **13**, 407–417 (1987).
- Swan, J. A. & Pearce, J. M. The orienting response as an index of stimulus associability in rats. *J. Exp. Psychol. Anim. Behav. Processes* **14**, 292–301 (1988).
- Wilson, P. N., Boumphrey, P. & Pearce, J. M. Restoration of the orienting response to a light by a change in its predictive accuracy. *Q. J. Exp. Psychol.* **44B**, 17–36 (1992).
- Miller, R. R. & Matzel, L. D. in *Contemporary Learning Theories: Pavlovian Conditioning and the Status of Traditional Learning Theory* (eds., Klein, S. B. & Mowrer, R. R.) 61–84 (Erlbaum, Hillsdale, New Jersey, 1989).
- Gallistel, C. R. & Gibbon, J. Time, rate, and conditioning. *Psychol. Rev.* **107**, 289–344 (2000).
- Neal, R. M. *Bayesian Learning for Neural Networks* (Springer, New York, 1996).
- Moran J. & Desimone R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
- Lu, Z. L. & Doshier, B. A. External noise distinguishes attention mechanisms. *Vision Res.* **38**, 1183–1198 (1998).
- McAdams, C. J. & Maunsell, J. H. R. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
- Treue, S. & Martinez Trujillo, J. C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575–579 (1999).
- Cohen, J. D. & Servan-Schreiber, D. A theory of dopamine function and its role in cognitive deficits in schizophrenia. *Schizophrenia Bull.* **19**, 85–104 (1993).
- Han, J.-S., Holland, P. C. & Gallagher, M. Disconnection of the amygdala

- central nucleus and substantia innominata/nucleus basalis disrupts increments in conditioned stimulus processing in rats. *Behav. Neurosci.* 113, 143–151 (1999).
33. Corwin, J. V. & Reep, R. L. Rodent posterior parietal cortex as a component of a cortical network mediating directed spatial attention. *Psychobiology* 26, 87–102 (1999).
 34. Colby, C. L. & Goldberg, M. E. Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 22, 319–349 (1999).
 35. Mesulam, M. M. Spatial attention and neglect: parietal, frontal and cingulate contributions to the mental representation and attentional targeting of salient extrapersonal events. *Phil. Trans. R. Soc. Lond. B Biol. Sci.* 354, 1325–1346 (1999).
 36. Lubow, R. E. *Latent Inhibition and Conditioned Attention Theory* (Cambridge Univ. Press, New York, 1989).
 37. Solomon, P. R. & Moore, J. W. Latent inhibition and stimulus generalization of the classically conditioned nictitating membrane response in rabbits (*Oryctolagus cuniculus*) following hippocampal ablation. *J. Comp. Physiol. Psychol.* 89, 1192–1203 (1975).
 38. Baxter, M. G., Holland, P. C. & Gallagher, M. Disruption of decrements in conditioned stimulus processing by selective removal of hippocampal cholinergic input. *J. Neurosci.* 17, 5230–5236 (1997).
 39. Weiner, I. Neural substrates of latent inhibition: The switching model. *Psychol. Bull.* 108, 442–461 (1990).
 40. Eichenbaum, H. The hippocampus: The shock of the new. *Curr. Biol.* 9, R482–R484 (1999).
 41. Houk, J. C., Davis, J. L. & Beiser, D. G. (eds.) *Models of Information Processing in the Basal Ganglia* (MIT Press, Cambridge, Massachusetts, 1995).
 42. Wickens, J. & Kötter, R. in *Models of Information Processing in the Basal Ganglia* (eds., Houk, J. C., Davis, J. L. & Beiser, D. G.) 187–214 (MIT Press, Cambridge, Massachusetts, 1995).
 43. Kropotov J. D. & Etlinger S. C. Selection of actions in the basal ganglia-thalamocortical circuits: review and model. *Int. J. Psychophysiol.* 31, 197–217 (1999).
 44. Doya, K. What are the computations of the cerebellum, the basal ganglia and the cerebral cortex? *Neural Networks* 12, 961–974 (1999).
 45. Hatfield, T., Han, J.-S., Conley, M., Gallagher, M. & Holland, P. C. Neurotoxic lesions of basolateral, but not central, amygdala interfere with Pavlovian second-order conditioning and reinforcer devaluation effects. *J. Neurosci.* 16, 5256–5265 (1996).
 46. Schoenbaum, G., Chiba, A. A. & Gallagher, M. Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J. Neurosci.* 19, 1876–1884 (1999).
 47. Schultz, W., Tremblay, L. & Hollerman, J. R. Reward processing in primate orbitofrontal cortex and basal ganglia. *Cereb. Cortex* 10, 272–283 (2000).
 48. Ballard, D. H. Animate vision. *Artificial Intelligence* 48, 57–86 (1991).
 49. Schultz, W. & Dickinson, A. Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* 23, 473–500 (2000).
 50. Rescorla, R. A. & Wagner, A. R. in *Classical Conditioning II: Current Research and Theory* (eds., Black, A. H. & Prokasy, W. F.), 64–69 (Appleton-Century-Crofts, New York, 1972).
 51. Sutton, R. S. Learning to predict by the methods of temporal difference. *Machine Learning* 3, 9–44 (1988).
 52. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* 275, 1593–1599 (1997).
 53. Bordley, R. F. The combination of forecasts: a Bayesian approach. *J. Operational Res. Soc.* 33, 171–174 (1982).
 54. Lindley, D. V. in *Bayesian Statistics 2* (eds., Bernardo, J. M., DeGroot, M. M., Lindley, D. V. & Smith, A. F. M.) 375–390 (North Holland, Amsterdam, 1985).
 55. Jacobs, R. A., Jordan, M. I. & Barto, A. G. Task decomposition through competition in a modular connectionist architecture: the what and where vision tasks. *Cogn. Sci.* 15, 219–250 (1991).