PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

The Conditioning Connection

I. P. L. McLaren and A. Dickinson

Phil. Trans. R. Soc. Lond. B 1990 **329**, 179-186
doi: 10.1098/rstb.1990.0163

The conditioning connection

I. P. L. McLAREN AND A. DICKINSON

Department of Experimental Psychology, University of Cambridge, Downing Street, Cambridge, CB2 3EB, U.K.

SUMMARY

In 1948 Konorski argued that conditioning reflects the strengthening of a connection between elements representing the signal and the reinforcer, as a result of the coincidence of activity in the signal element with a rise in activity in the reinforcer element. This Konorskian process represents one way of implementing an error-correcting learning rule and thus, unlike a simple Hebbian process, anticipates selective conditioning, such as that observed in Kamin's blocking procedure. However, the Konorskian process, in common with other error-correcting learning rules, fails to explain why blocking is attentuated by 'surprising' changes in reinforcement conditions that should not augment activity in the reinforcer element. Rather, the reinforcer-specificity of such unblocking suggests the operation of an associability process by which stored information about the past predictive history of the signal is expressed at the connection between signal and reinforcer elements to modulate changes in the connection weight.

1. INTRODUCTION

The idea of associationism, namely that the learning consists of the formation of connections between mental entities, is an attractive one, not least because it attempts to explain learning at a psychological level in terms of processes that are in general terms compatible with our understanding of neural functioning. Indeed, so strong has been the assumed compatibility that often the psychological level of explanation disappears altogether when the nature of neural processes is inferred directly from behavioural observation alone. Pavlov (1927), for example, treated his conditioned response almost as a behavioural 'electrode' recording processes of 'irradiation' and 'concentration' in the cerebral cortex. The dangers in such a strategy are obvious; operating, as he was, outside the Sherringtonian concept of the nervous system, Pavlov was led, purely on the basis of behavioural data, to explanations couched in terms of fictitious neural processes. Konorski (1948) attempted to rectify this particular deficit when he carefully reinterpreted the phenomena of Pavlovian conditioning in terms of a Sherringtonian nervous system and, as is well known, Hebb (1949) subsequently extended this type of theorising to a variety of psychological functions.

Once the dangers of creating neural fictions on the basis of behavioural observations alone is recognized, connectionism still remains an attractive theory of learning on the psychological level because of its potential for a reduction to neural mechanisms. This potential depends, however, upon constraining the psychological theory by some very general principles derived from our conception of neural architecture and functioning. Apart from the associative assumption that learning consists of the strengthening (or weakening) of connections between 'representational' elements, both Konorski (1948) and Hebb (1949) appeared to impose one further very general constraint, namely, that changes in connection strength or weights are solely a function of the concurrent activity in the two elements of the connection. This we shall refer to as the local activity constraint, a constraint that would appear to be endorsed by many contemporary neural connectionists (Levy & Desmond 1985). Put more precisely, the constraint is in two parts: first, that changes in the strength of a connection between two representational elements should depend in some fashion on the activity of those elements, and secondly, that these activities are in turn determined in some simple way by the stimuli they represent. In terms familiar to the learning theorists, changes in the strength of the connection between elements representing the signal and reinforcer should be determined by the activities of those elements (or some function of them), and those activities should reflect solely the properties of the stimuli involved. Our purpose in the present paper is to examine whether a psychological theory of learning operating within this constraint can explain the basic properties of simple associative learning and, if not, in what ways this constraint has to be relaxed.

The form of associative learning against which we have chosen to assess this constraint is that addressed by Konorski (1948): simple conditioning in which learning about an association between a signal and reinforcer is indexed by observing the acquisition of a new response to the signal. Our reason for this choice is twofold. First, we believe that simple conditioning is a form of learning that directly reflects the principles governing the formation of individual connections. There is no doubt that the renaissance of associationism in psychology is primarily because of demonstrations that parallel-distributed architectures (Rumelhart &

McClelland 1986) allow this view of learning to embrace phenomena that have always been problematic for associative accounts, such as generalization, perceptual learning and concept formation (McLaren *et al.* 1989). However, the associationist explanations of these psychological functions, by their very nature, assume that the phenomena reflect the interaction of many modifiable connections so that the principles governing individual weight changes may be masked by the emergent properties of the network. The conditioning procedure, by using relatively simple stimuli and associative relationships, is assumed to unmask the basic principles of connection by minimizing the role of these emergent properties. Secondly, there is a realistic prospect of rapprochement between psychological and neural connectionism in the case of simple conditioning (see contributions by Thompson and Carew (this symposium)).

2. THE HEBBIAN RULE

Although it is Hebbian rather than Konorskian learning rules that figure largely in contemporary neurobiologically motivated connectionist thinking (see, for example, Levy & Desmond 1985), we have known for some time that a simple Hebbian rule is insufficient for simple conditioning. A Hebbian rule assumes that conjoint activity in the input and receptor elements, the signal and reinforcer 'representations' respectively, is sufficient to strengthen the connection weight between the elements, (figure 1*a*). Ever since Kamin's (1969) well-known demonstration of 'blocking', the inadequacy of the Hebbian process has been apparent.

We shall show this blocking effect by an appetitive conditioning study from our own laboratory (see table 1). In the first stage one group of rats, group L/0-0, is trained to press a lever in the presence of a light [L] for a food reward that is delivered at the end of this signal (see Dickinson & Mackintosh (1979) for details of the training procedure). When responding is established to the light, a second signal, a clicker [C], is presented in simultaneous compound with the light and the reinforced training continued for a number of trials. The question of interest is whether the animals learn normally about the association between the clicker and the food during this second stage. A simple Hebbian theory predicts that they should; during compound training, activity in the elements that represent the clicker and food should overlap in time on each trial, thus producing increments in their connection weight. As figure 2 shows, however, the presentation of the clicker by itself after compound training produced little or no increase in responding, suggesting that the animals had learned little about the asssociation between the signal and reinforcer during this stage of training.

That the failure to observe conditioning to the clicker was not simply a result of the ineffectiveness of the clicker as a signal during compound training is shown by a second, control group, group D/0-0. These animals received the same training as a group L/0-0 except for the fact that the signal for food in the first

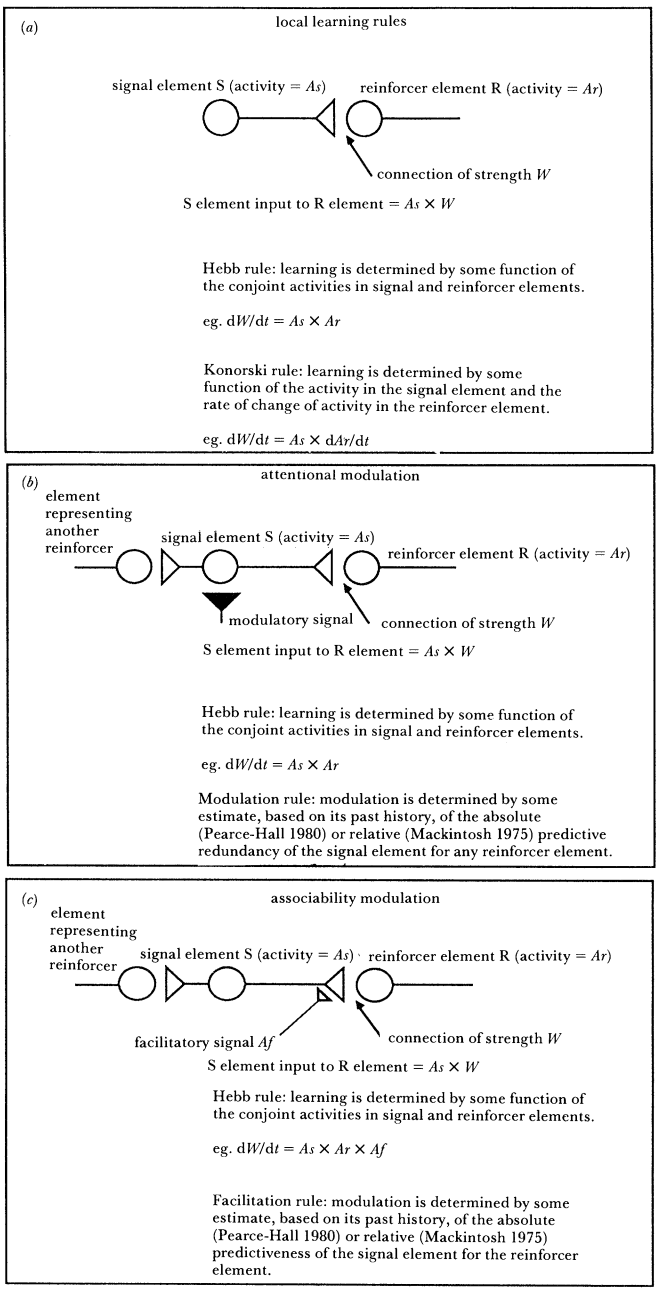

Figure 1. (*a*) Examples of Hebbian and Konorskian learning rules. (*b*) How some kind of attentional modulation might be implemented. Note that modulation in this case depends on the predictive redundancy of the signal(s) for all reinforcers, and that it affects all learning involving a given signal. (*c*) A suitable architecture for the implementation of associability modulation. Now the modulation is reinforcer specific, both in its determination and application.

Table 1. *Design of blocking experiment*
(L, light; D, darkness; C, clicker; F, food presentation.)

group	stage 1 (15 trials)	stage 2 (10 trials)
L/0-0	L- > F	CL- > F
D/0-0	D- > F	CL- > F
L/0-F	L- > F	CL- > F-F
L/F-F	L- > F-F	CL- > F-F
L/F-0	L- > F-F	CL- > F

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

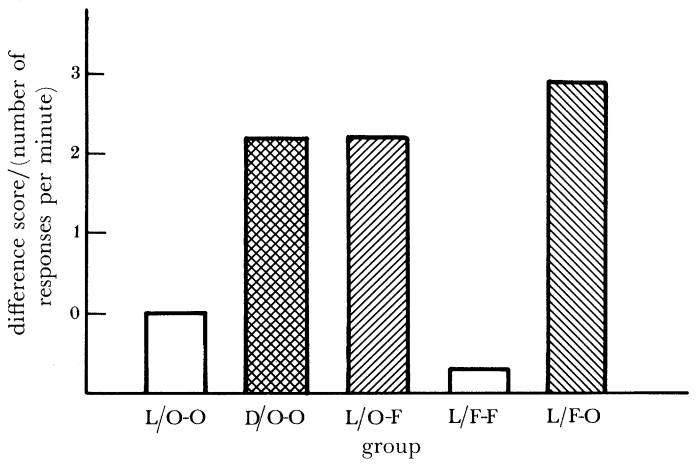

Figure 2. Mean difference scores for the various groups in the blocking experiment. The difference scores reflect the difference between the number of responses per minute to the clicker on test and in an equivalent pre-signal period. F tests using the overall error term revealed that the following planned comparisons were significant at $p < 0.05$: L/0-0 against D/0-0; L/0-0 against L/0-F; L/F-F against L/0-F; L/0-0 against L/F-0.

stage was darkness [D] produced by switching off the house light in the chamber rather than switching on an additional light [L]. The fact that the presentation of the clicker alone produced a significant increase in responding on test in group D/0-0 shows that the compound training could support learning about the signal-food relation. A comparison of the test performance to the clicker in the two groups suggests that the training with the light in the first stage 'blocked' learning about the clicker during compound conditioning in the second stage for group L/0-0.

This blocking effect represents a general property of simple conditioning, namely that the learning process discriminates against redundant signals. Given that the food reward is adequately predicted by the light in group L/0-0, the presence of the clicker during compound training provides no further information about the occurrence of the food. The failure of Hebbian processes, to explain the sensitivity of conditioning to signal redundancy is a major limitation, which cannot be avoided by introducing extra rules that apply when either signal or reinforcer are presented on their own (e.g. the homosynaptic and heterosynaptic rules discussed by Willshaw & Dayan (1990)), so that the system is sensitive to event correlation rather than simple pairing.

3. THE KONORSKIAN RULE

One response to Kamin's results with the blocking procedure was the development of the Rescorla–Wagner model for conditioning. This employed an error-correcting rule, in which the presence of a reinforcer sets a target value for the net associative strength of all signals present to that reinforcer. Learning is then driven by the discrepancy between actual and target associative strength. Blocking is explained by the pre-trained signal acquiring sufficient associative strength for the reinforcer to meet, or nearly meet, the target requirement set by presentation of the reinforcer, leaving little associative strength to accrue to the second signal when presented in compound with the first. This type of error-correcting rule is impossible to instantiate with signal and reinforcer elements employing a pure Hebbian rule. Either a third component carrying the error signal has to be introduced which violates the local activity assumption, or intra-element processes must be postulated to perform the necessary computations†. The first option, involving the construction of a neural assembly to serve as an error-correcting processing element, is discussed in McLaren (1989) and Donegan *et al.* (1989). The second alternative is considered here. Although a number of theorists have simply assumed that weight changes are governed by an error-correcting rule (see, for example, McClelland & Rumelhart 1985), as far as we know, Konorski (1948) was the first to propose an intra-element process that can implement such a rule. Surprisingly, Konorski's (1948) connectionist theory has received little or no attention, even though he gave a much more precise specification of the processes controlling changes in connection weights than did Hebb (1949). And, in fact, his learning rule anticipates blocking even though he himself was not aware of the effect. Konorski essentially argued that an increment in the connection weight was brought about by the coincidence of activity in the input element with rising activity in the receptor element, an idea that has been developed independently by Sutton & Barto (1981). This can be implemented by postulating two processes in the element representing the reinforcer. One, the *P* process, drives ordinary Hebbian learning, the other, *N* process, produces anti-Hebbian learning (i.e. a decrease in connection strength that is a function of the conjoint activity in signal and reinforcer elements). Both *P* and *N* depend on the activity of the element, but they have different time-constants in responding to changes in activation, with *P* reacting more quickly than *N*.

† Mitchison (1989) has suggested that another solution is to separate the necessary computations in time. The reader will see that our implementation of Konorski's rule does in fact make use of this strategy, in that the Hebbian and anti-Hebbian processes follow different timecourses.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

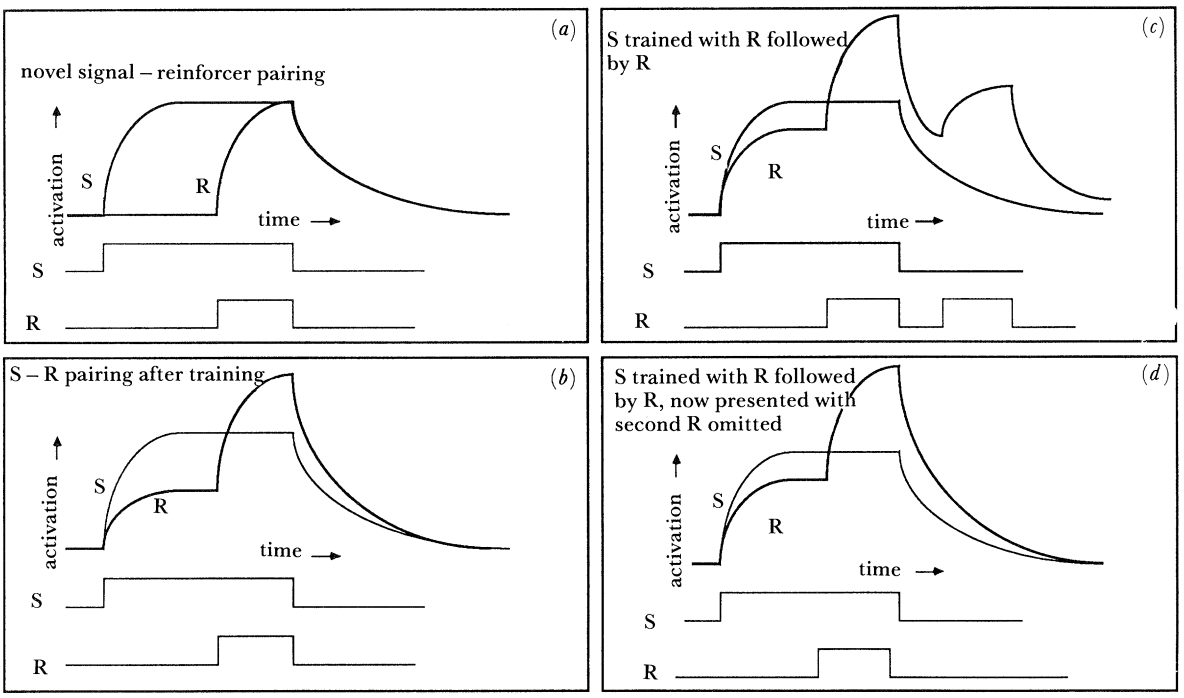

Figure 3. (*a*) The activation profile over time for signal (S) and reinforcer (R) elements that have not previously been paired. The lower portion of the figure depicts the onset and offset of the signal and reinforcer. (*b*) The final activation profile after many such pairings. Now learning and unlearning are in balance over a complete cycle. (*c*) The effect of adding another reinforcer presentation after the first. The result is further learning, until the net learning in a cycle is zero once again (as shown in this figure). (*d*) The post-trial reinforcer has been removed, leaving more unlearning than learning during a cycle.

Hence, when the activation of the reinforcer element is rising, $P > N$ and the connection is strengthened, whereas when the reinforcer activation is falling $N > P$ and the connection is weakened. If the activation is constant, then P and N settle to equal equilibrium values and the net learning is zero. In all the cases considered so far it is assumed that the signal element is active, and in effect the rule is that learning depends on the product of signal activation and some function of the rate of change of reinforcer activity (see figure 1). This function is chosen so that simultaneous activation of signal and reinforcer elements (which would occur when extinguishing a previously conditioned signal) results in a decrement in connection strength. A positive increment is contingent on the signal preceding the reinforcer by some time, there being a temporal 'window' surrounding an optimal signal-reinforcer interval.

If the appropriate assumptions are made, this rule can be shown to be of the error-correcting type. Simulations of Konorski's theory have shown that it can explain blocking, a fact that is also clear from an intuitive application of the theory. Figure 3*a* shows the state of affairs when a new signal and reinforcer are paired; each comes on independently of the other, in the sense that the activations in their respective representational elements are determined only by the signal or reinforcer in question. Learning occurs because signal element activity is paired with an increase in reinforcer element activity (at the onset of this stimulus). By contrast, unlearning occurs when signal element activity is paired with a fall in reinforcer element activity, which takes place at reinforcer offset in figure 3*a*†. Note, however, that there is a bias in favour of learning because signal element activity is greater when activity in the reinforcer element is rising rather than falling. The activity profiles after learning are shown in figure 3*b*. On each trial there is some extra learning due to the rise in reinforcer element activity as the signal element is activated, but this is counteracted by the unlearning that now occurs at reinforcer offset, because of the very rapid fall in reinforcer activity from a high level. The net effect of such a learning trial will tend towards zero, at which point a stable connection has been formed; indeed, we can take figure 3*b* to depict just such a state of affairs.

Now we can see why blocking occurs. A trained signal, such as the light presented to group L/0-0 during compoung conditioning, influences the activity in the reinforcer element such that the net learning over a cycle for that signal–reinforcer relation is zero. Pairing another signal such as the clicker in the same temporal relationship to the reinforcer as the light simply exposes it to the same zero change cycle and little learning occurs. This account shares some features with the model proposed by Hawkins (1989), in that the temporal relation between signal and reinforcer plays a central role in learning, whilst differing from the latter in that no complex arrangement of elements is required.

† Konorski (1948) himself argued that the pairing of input element activity a fall in receptor element activity led to the development of an inhibitory connection rather that the weakening of an excitatory connection. For our present purpose, however, we can assume that this process brings about unlearning.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

Konorski's theory also predicts conditions under which blocking could be prevented. Suppose that a second, post-trial food presentation is given sometime after the first on each compound trial, as shown in figure 3*c*. If the delay to the second food presentation is sufficient to allow some decay of activity in the reinforcer element, the second presentation should produce another rise in activity which in turn should support a further weight increase, assuming of course that there is still some residual activity in the signal element. Kamin (1969) himself, in fact, reported that the delivery of a second, unexpected reinforcer shortly after the first reinforcer on compound trials restored conditioning to the normally blocked signal. This unblocking effect was replicated in our study in group L/0-F. These animals initially received training in which the light signal was paired with a single food presentation in the first stage. During compound training, however, each light-clicker compound signalled not only the first reinforcer but also a second food presentation 8 s later (see table 1). In our study the presentation of this second reinforcer completely restored conditioning to the clicker on test relative to the control group, group D/0-0 (see figure 2). Of course, if the double reinforcer is presented throughout both stages of training, some blocking should be reinstated. Pairing the light with a double reinforcer in the first stage should establish a stronger connection weight between their respective elements than if only the first food presentation is presented. This means the presence of the light on compound trials should induce a higher level of activity in the reinforcer element at the time of both food presentations, as shown in figure 3*c*. Again a zero change cycle will be achieved, and blocking should occur. This restoration of blocking was evident in group L/F-F which received the double reinforcer during both stages of training (see table 1); there was no evidence of any conditioning to the clicker on test in this group (see figure 2).

The pattern of conditioning observed in this blocking study accords with the general claim that little learning occurs to a signal that is a redundant predictor. Moreover, the Konorskian learning rule appears to account for the effects of signal predictiveness within the local activity constraint. There is, however, at least one condition that dissociates predictive redundancy and Konorski's rule, a condition that was implemented in the last group of our study (see table 1). The animals in group L/F-0 received a double reinforcer paired with the light during the first stage but only a single food presentation in association with the clicker-light compound during the second stage. In this group the clicker is not a redundant predictor in that it signals the omission of the second reinforcer, and yet this is not a change in reinforcement conditions that should bring about conditioning to the clicker according to Konorski; if anything it should bring about unlearning to the clicker because activity in the signal element will be paired with an even greater fall in activity in the reinforcer element (see figure 3*d*). The fact that group L/F-0, if anything, shows the most conditioning to the clicker (see figure 2) clearly presents a considerable challenge to the Konorskian rule.

The subtle effects of predictive redundancy on conditioning are not a peculiar artefact of our particular procedure; the attenuation of blocking by unexpected reinforcer addition and omission has been replicated in different conditioning procedures and laboratories (see, for example, Dickinson *et al.* 1976; Holland 1984, 1988). And, as a central feature of simple conditioning, sensitivity to predictive redundancy must be addressed by any connectionist model that has pretensions to generality. The failure of Konorski's theory, and indeed simple error correcting theories in general (see, for example, Rescorla & Wagner 1972; Sutton & Barto 1981; McClelland & Rumelhart 1985), in the face of unblocking by reinforcer omission suggests that the local activity constraint must be relaxed if we are to provide an account of this effect.

4. ATTENTIONAL PROCESSES

Once we have relaxed the local activity constraint, other ways of making a connectionist system sensitive to predictive redundancy become viable. An obvious possibility is to implement some form of attentional system to ensure that the activity in a signal element is determined by the signal's predictive redundancy, violating the second part of the constraint. For example, both Mackintosh (1975) and Pearce & Hall (1980) suggest a basic associative learning rule that could be implemented by a simple Hebbian process for weight changes. They also argue, however, that the degree to which a signal will enter into association with a reinforcer depends upon whether or not the subject attends to the signal. Attention to a signal is in turn governed by its past predictive redundancy; if the signal has been redundant in the past, it will receive little attention, resulting in a failure to learn about its association with a reinforcer on a current trial.

The advantage of such an attentional theory is its ability to explain the effect of predictive redundancy observed in the blocking procedure. In the groups which exhibit blocking, groups L/0-0 and L/F-F, the conditions of reinforcement on the compound trials are exactly those anticipated on the basis of the training to the light in the first stage. Thus the animals should detect the redundancy of the clicker on the initial compound trials and rapidly cease to attend to it with the consequence that they learn little about its relationship to the reinforcer. Of course, the change in reinforcement conditions produced by the addition and omission of the second reinforcer on each compound trial in groups L/0-F and L/F-0, respectively, should maintain attention to the tone and hence learning to it.

Given the apparent concordance between the predictions of attentional theory and the conditions for unblocking, the next step is to embody an attentional process within a connectionist architecture. An obvious way of doing this, which is in keeping with the psychological processes identified by Mackintosh (1975) and Pearce & Hall (1980), would be to construct a circuit that modulates the activity in the signal element in a way that reflects the predictive redundancy of the signal. A Hebbian or Konorskian

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

rule could then be used to change the weight of the connection between the signal and reinforcer elements. Figure 1*b* shows an example of this approach.

The attentional theory argues that an agent will learn about the relation between a signal and a particular, target reinforcer as a result of their pairing as long as the signal has not been a redundant predictor. Note, however, that this theory does not require the signal to have been informative about the target reinforcer itself; it is sufficient that the signal has been informative about the occurrence of some important event, for then it should continue to receive attention and thus be associated with the target reinforcer. This is true even if the signal has been a redundant predictor of the reinforcer itself.

Dickinson & Mackintosh (1979) investigated whether it is the general predictive redundancy of a signal that determines learning within the context of an unblocking study similar to that we have already described. Four groups replicated the unblocking effect by both the surprising addition and omission of the second food reinforcer. More important in the present context, however, is the conditioning to the clicker in a second quartet of groups. These rats received exactly the same training as the first set except for the fact that the second food presentations were replaced by the delivery of a mild foot shock. Thus for group L/0-S the light was associated with a single food presentation in the first stage. Compound trials with the light and clicker also terminated with a food presentation but in addition a shock was presented 8 s after the end of each trial. In this group the clicker is redundant as a predictor of the food reinforcer, while being informative about the occurrence of the shock. Hence, if learning is determined by predictive redundancy in general, we should have expected these animals to maintain their attention to the clicker during compound training and thus have learned about its association with the food reinforcer. This did not appear to happen; presentation of the clicker alone on test produced little elevation of responding and certainly no more than that observed in group L/S-S in which the clicker was redundant with respect to both the food and shock. In this group the light was established as a predictor of both the food and the shock in the first stage.

There is, of course, an entirely trivial explanation of the failure to find unblocking by the unexpected or surprising addition of the shock. Although Dickinson & Mackintosh (1979) hoped that delaying the shock for 8 s after each compound trial would minimize any direct aversive conditioning, it is possible that the low level of responding to the clicker reflected aversive conditioning to the clicker rather than the blocking of appetitive conditioning to this signal by the light. The results of the final two groups, however, render this possibility unlikely. The light was established as a predictor of both the food and shock reinforcers in the first stage for group L/S-0 before the animals received training with the clicker-light compound paired with the food alone. In this group the clicker was not a redundant predictor because it signalled the absence of the shock in fashion analogous to the prediction of the absence of the food in group L/F-0. Nor could responding be suppressed by direct aversive conditioning to the clicker in group L/S-0 as the shock was never presented in association with the clicker. Thus if attention to a signal is determined by its general predictive redundancy, conditioning should have been observed to the clicker in this group, but it was not.

In conclusion, rendering the clicker informative by making it predictive of either the addition or omission of the shock did not prevent the light blocking appetitive conditioning to the clicker. The minimal level of responding to the clicker in groups L/0-S and L/S-0 corresponds to that observed for the control group, group L/0-0, that simply received the appetitive blocking procedure. This pattern of results stands in marked contrast to that observed when predictive redundancy was manipulated by the addition or omission of a second food reinforcer.

5. ASSOCIABILITY PROCESSES

These results showed that for a pairing of a signal and reinforcer to bring about learning of their association, it is not sufficient that the signal is a generally informative stimulus. Rather it must be informative relative to other signals present about the occurrence of the reinforcer that supports conditioning. This idea can be captured by the concept of reinforcer-specific associability. The associability of a signal with a particular reinforcer is determined by whether or not it is a redundant predictor of that reinforcer. The signal will remain associable with a reinforcer only as long as it is a good predictor of the occurrence (or non-occurrence) of that reinforcer. This claim predicts that if we were to reverse the role of the food and shock in our blocking experiment to generate an aversive rather than appetitive conditioning procedure, the reverse pattern of results should be observed. Thus we could study the effect of prior aversive condition to the light on its ability to block conditioning to the clicker when a clicker-light compound is paired with a shock reinforcer. Now it should be the unexpected addition and omission of a second shock rather than food presentations and omissions that should produce unblocking and restore conditioning to the clicker. Without going into details, this is just what Dickinson & Mackintosh (1979) observed.

An implementation of the requirement for reinforcer-specific associability is shown in figure 1*c*. The activities of the signal and reinforcer elements are set in a simple fashion by presentation of the signal and reinforcer, but now it is the learning at a particular connection that is modulated by computation of predictive redundancy. Specifically, it is the extent to which the signal has been predictive of the reinforcer in the past that determines the modulation at this connection. We can imagine the reinforcer element assessing the extent to which it is predicted, and then delivering the results of this assessment to the incoming connections to it, so that the modulation can be gradually altered rather than instantaneously set. The modulatory influence would only be effective when the signal element was active, in which case it would

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

determine the rate at which the connection weight could change. This model instantiates a reinforcer-specific learning parameter, for which we now reserve the term associability, leaving open the question of whether or not changes in associability affect the expression of an association in performance. To see this, consider that if the modulatory signal were to influence the activity of the signal element directly, then its effect would no longer be reinforcer-specific, which is the reason why the modulation is shown as being accomplished via a connection onto the plastic connection in question. The modulatory connection might have a facilitatory effect on learning, such as is found in Aplysia (Hawkins 1989), in which case its effect on performance might be weak or non-existent as assumed in McLaren (1989). If, however, associability is to be expressed in performance as well as learning, then the modulation must affect activity local to the connection. This modulation would be of an inhibitory type, so that there would be no observable effect unless the signal element were active; otherwise the computations necessary are similar to those in the previous case. Clearly, the modulation need not affect the activity of the signal element itself, thus avoiding the problems of the attentional approach while implementing an associability parameter whose value will also be felt in performance.

Whether or not associability affects the expression of associations between the relevant signal and reinforcer representations is an empirical question that has not received a great deal of attention. Mackintosh (1975) suggests that it may do so, and Hall *et al.* (1977) present evidence that can be construed in support of this claim. They showed that a weak signal pre-trained to a reinforcer and then trained in compound with a strong signal to the same reinforcer lost, to a large extent, its ability to evoke a response when tested in isolation. If it is assumed that the strong signal gained sufficient associative strength over a few trials to become a better predictor of the reinforcer than the weak signal, then the associability of the weak signal should decline, decreasing the strength of the response controlled by this stimulus.

The reader should also note that the modulatory influence is not instantaneously set; rather it is shifted up and down in a gradual fashion whenever the signal element is active and the necessary computations are performed at the reinforcer element. The associability parameter is unchanged during periods when the relevant elements are inactive. We think of the associability as starting relatively high and having considerable inertia, so that even in the case of a redundant stimulus the decline in associability is smooth rather than sudden. Changes in associability reflect the history of the signal–reinforcer relation up to that point, when considered in conjunction with the relation between other signals and that reinforcer. It is also worth pointing out that the computations determining associability probably embrace a longer time interval than those determining changes in connection strength *per se* (see Holland 1988).

Our argument is that the behavioural data constrain the connectionist architecture to a certain minimum level of complexity. The concept of associability requires fairly sophisticated computations to be performed and then applied at the site of the connection between the signal element (or some intermediary) and the element representing the reinforcer. The implication is that we should expect to find more than just the processes directly involved in changes in connection weight at this locus. As presented, our implementation of associability clearly violates the local activity constraint; extra machinery is invoked, either to influence learning directly or to modulate activity at the locus of the connection. This represents an infringement of the first clause of the constraint, as signal element activity *per se* is able to reflect the signalling stimulus' properties in a simple way, but signal and reinforcer element activity do not totally determine learning at the connection. Imagine, however, that by having local modulation of connection activity and sufficient processes operating intra-element to implement the computations for associability, the architecture shown in figure 3*c* could be reduced to one element connecting with another. Does this reduction serve to rescue the local activity constraint? The answer is no, because the activity involved in learning is not simply signal and reinforcer activities, but rather some locally modulated activity at the connection. Again, if learning were directly influenced by some parameter other than activity inside the signal element, then this still violates the constraint. Figure 1*c* shows this extra factor explicitly for ease of exposition, but it must be represented in some form, and learning cannot simply be a function of signal and reinforcer element activity. This makes good computational sense. Surely the activities of signal and reinforcer elements should be reserved to represent properties of the signal and reinforcer, not the relationship between them in the context of other stimuli. That is more appropriately done at the locus of the connection between signal and reinforcer. Such an arrangement is also easier to implement, because all the information necessary to compute associability can be available locally at the connection. So we see that while learning can still be a local process, which is computationally desirable, this process cannot reflect just a simple function of the activity in the signal and reinforcer elements.

Whatever the merits of the various possible implementations of an associability mechanism, it is clear that some such process is necessary if we are to explain selective learning in simple conditioning. An error-correcting learning rule is not sufficient in itself. Whether or not we require some form of error correction, in addition to an associability mechanism, at present, remains unclear. If we do, Konorski's learning rule should be considered as a serious candidate for the implementation of such a process.

REFERENCES

Dickinson, A., Hall, G. & Mackintosh, N. J. 1976 Surprise and the attentuation of blocking. *J. exp. Psychol.*: *Anim. Behav. Proc.* **2**, 313–322.

Dickinson, A. & Mackintosh, N. J. 1979 Reinforcer speci-

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES

ficity in the enhancement of conditioning by posttrial surprise. *J. exp. Psychol.*: *Anim. Behav. Proc.* **5**, 162–177.

Donegan, N. H., Gluck, M. A. & Thompson, R. F. 1989 Integrating behavioural and biological models of classical conditioning. In *The psychology of learning and motivation: computational models of learning in simple neural systems* (ed. G. H. Bower). London: Academic Press.

Hall, G., Mackintosh, N. J., Goodall, G. & Martello, M. 1977 Loss of control by a less valid or by a less salient stimulus compounded with a better predictor of reinforcement. *Learn. Motiv.* **8**, 145–158.

Hawkins, R. D. 1989 A biologically based computational model for several simple forms of learning. In *The psychology of learning and motivation: computational models of learning in simple neural systems* (ed. G. H. Bower). London: Academic Press.

Hebb, D. O. 1949 *The organisation of behaviour*. New York: Wiley & Sons.

Holland, P. C. 1984 Unblocking in Pavlovian appettitive conditioning. *J. exp. Psychol.*: *Anim. Behav. Proc.* **10**, 476–497.

Holland, P. C. 1988 Excitation and inhibition in unblocking. *J. exp. Psychol.*: *Anim. Behav. Proc.* **14**, 261–279.

Kamin, L. J. 1968 'Attention-like' processes in classical conditioning. In *Miami symposium on the prediction of behaviour: aversive stimulation* (ed. M. R. Jones), pp. 9–33. University of Miami Press.

Konorski, J. 1948 *Conditioned reflexes and neuron organisation*. Cambridge University Press.

Levy, W. B. & Desmond, N. L. 1985 The rules of elemental synaptic plasticity. In *Synaptic modification, neuron* (ed. W. B. Levy, J. A. Anderson & S. Lehmkuhle), pp. 105–122. New Jersey: Lawrence Erlbaum Associates.

Mackintosh, N. J. 1975 A theory of attention: variations in the associability of stimuli with reinforcement. *Psychol. Rev.* **82**, 276–298.

McClelland, J. L. & Rumelhart, D. E. 1985 Distributed memory and the representation of general and specific information. *J. exp. Psychol.*: *Gen.* **114**, 159–188.

McLaren, I. P. L. 1989 The computational unit as an assembly of neurones: an implementation of an error correcting learning algorithm. In *The computing neuron* (ed. R. Durbin, C. Miall & G. Mitchison). Amsterdam: Addison Wesley.

McLaren, I. P. L., Kaye, H. & Mackintosh, N. J. 1989 An associative theory of the representation of stimuli: applications to perceptual learning and latent inhibition. In *Parallel distributed processing – implications for psychology and neurobiology* (ed. R. G. M. Morris). Oxford University Press.

Mitchison, G. 1989 Learning algorithms and networks of neurons. In *The computing neuron* (ed. R. Durbin, C. Miall & G. Mitchison). Amsterdam: Addison Wesley.

Pavlov, I. P. 1927 *Conditioned reflexes*. Oxford University Press.

Pearce, J. M. & Hall, G. 1980 A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* **87**, 532–552.

Rescorla, R. A. & Wagner, A. R. 1972 A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: current research and theory* (ed. A. H. Black & W. F. Prokasy), pp. 64–99. New York: Appleton–Century–Crofts.

Rumelhart, D. E. & McClelland, J. L. 1986 *Parallel distributed processing*, vol. I. Massachusetts: Bradford Books.

Sutton, R. S. & Barto, A. G. 1981 Toward a modern theory of adaptive networks: expectation and prediction. *Psychol. Rev.* **88**, 135–170.

Willshaw, D. & Dayan, P. 1990 Optimal plasticity from matrix memories: what goes up must come down. *Neural Comput.* (In the press.)

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B BIOLOGICAL SCIENCES