

Studies in animal learning

Tony Ludlow
01420 83922
tony.ludlow@modelresearch.cix.co.uk

February 13, 2003

“that process which manifests itself by adaptive changes in individual behaviour as a result of experience.” *Thorpe (1963)*

Manning and Dawkins (1992), and their more recent edition, both state that the old approach to learning has “been replaced by a much more biologically based approach. ... learning is now seen as a way in which animals attempt to identify key aspects of a fluctuating environment: to detect its regularities and ignore the distracting ‘noise’ which is not important for them.”

My aim in these lectures is to discuss how animals pick out the relevant detail and much of the material is drawn from a classic book by *Dickinson (1980)* recommended by Manning and Dawkins. The ideas in the book are important but it is a hard read. I have tried to extract the best bits and place them in an intelligible framework. That means I have had to skip many other aspects of learning and memory but these are covered by *Manning and Dawkins (1992)* and in other introductory texts. Most of these commend the modern approach but make no attempt to explain it. An exception is *Pearce (1997, Chapter 3)* which gives a good account of theoretical developments since *Dickinson (1980)*.

1 The old classification of learning

Since Pavlov’s early work, learning studies have been a major branch of animal behaviour and a standard classification arose (eg *Thorpe, 1963*)

1. Habituation
2. Conditioned reflex Type I (= classical conditioning)
3. Conditioned reflex Type II (= Trial and error learning or instrumental conditioning)
4. Latent learning
5. Insight learning
6. Imprinting

1.1 Habituation

Habituation occurs when an animal gradually ceases to respond to a stimulus. The classic case is farm animals by a railway line. The first train sends them in a panic but, gradually, they cease to respond at all. It is an example of the precautionary principle: treat all sudden changes as dangerous unless they prove otherwise. Thus, caution is built into the genetic code but the animal learns the detail of what is and is not dangerous. Virtually all animal groups show habituation.

1.2 Associative learning

This covers type I and II conditioned reflexes. Broadly, classical conditioning, or type I, occurs when a stimulus such as a sound is presented at the same time as food. After a number of trials the animal starts to salivate when the sound is presented on its

own. The simplest measures of classical conditioning are the intensity or probability of the response.

Trial and error learning, or type II, occurs when an animal's own actions are consistently rewarded or punished. For example, pressing a lever is rewarded with food. The animal rapidly becomes a lever junky. Again, the simplest measures are the frequency or intensity of response.

By far the greatest effort has been devoted to studying associative learning and a huge theoretical framework has been built up with heated debate and disagreement. The ambition of many authors has been to identify "laws of learning" but this has failed.

1.3 Latent learning

Latent learning occurs when an animal learns about the world without producing a clear response. Increasing knowledge of local geography would be an example, but we shall come across others. One important thing about latent learning is that it has proved very hard to fit into theories of associative learning because they are designed to explain the animal's response and, *in latent learning, there is no response*. Dickinson (1980) has referred to this as "the problem of behavioural silence".

1.4 Insight learning

A rather loose term to indicate that the animal seems to be learning tasks that require integration and analysis of the environment.

1.5 Imprinting

A highly specialised type of learning when a very young animal at a critical stage in its development imprints on a parent or type of food. For example, Lorenz found that hand-reared ducklings treated him as a duck and disregarded real ducks.

Features of imprinting are that it is extremely rapid, it occurs at a short and precise stage in the young animal's life and is permanent. These features have attracted research attention over the last twenty years.

2 Trying to understand learning

Putting labels on different kinds of learning is not the same as understanding what is going on, and Dickinson (1980) produced the first integrated survey of a new approach to learning. Since it is impossible to cover the whole field in two lectures I have based the lectures on his approach because it covers some interesting questions and gives a flavour of the kind of evidence that is used.

Classical studies described behaviour in terms of a stimulus and a response. Both stimulus and response could be measured and so the physiologist or behaviourist using this approach could be objective. He was not trying to measure what was going on inside the animal's head, which would be too subjective.

The approach, however, focussed on the wrong kind of question. It found questions that could be answered by simple tests, rather than finding the best way to answer crucial questions. As a result, the behaviourists denied that latent learning existed because they could not explain it in terms of stimulus and response.

In science, it is always better to get an approximate answer to the important questions rather than get a very precise answer to trivial ones. Dickinson (1980) has tried to do that. What he, and other cognitive psychologists, want to know is how relevant information is selected; how it is represented and stored in the brain and how it is processed in behaviour.

2.1 Sensory preconditioning

Take the example of a rat exposed to trials in which a light is paired with an electric shock. The animal cannot escape and, if the shock is strong, the rat will probably crouch immobile when the light comes on.

What has the rat learned? For the behaviourist, it has learned a new response to the light: crouching or freezing.

A cognitive psychologist is likely to argue that the rat has learned that the light predicts the shock and the rat is therefore frightened by it.

The difference between the two approaches is shown if, before pairing the light and shock, we pair sound and light. Both are neutral stimuli and both produce no response. The behaviourist therefore argues that the rat has learned nothing.

However, if we then train the rat to expect a shock when the light comes on, we find that presenting the sound instead causes a fear response (Rizley and Rescorla, 1972). So the rat *had learned something*. It had learned that the sound predicted the light, but the learning could only be demonstrated when the light became a highly relevant stimulus.

So it is possible to prove that latent learning has occurred, and this is an example of ‘sensory preconditioning’. There are many other examples of behaviourally silent learning. This is also an example of ‘integration of knowledge’. The rat had learned about tone and light and was able to integrate this knowledge with its later knowledge of light and shock.

This simple example takes us a long way and raises the questions:

How was the knowledge represented in the rats brain?

How was it selected, how was it stored and how was it integrated?

Do rats store information in a procedural way (like Fortran or C) or in a declarative way (like Prolog or Lisp)?

These are far more interesting questions than those the behaviourists studied while they stuck to their self-imposed limit of stimulus and response. The aim of the lectures is to show how they can be tackled.

3 Types of association

The main contention in these studies is that animals possess mechanisms designed to detect and store information about causal relationships. The first step in this approach is to examine the kinds of association that can be formed between events.

3.1 The nature of the association

Associative learning revolves round what happens when two events (E1 and E2) are paired and Dickinson (1980) replaces the old names for associative learning by suggesting that we focus on four possible pairings between two events:

- E1 → E2
- E1 → no E2
- Action → E2
- Action → no E2

Classical conditioning studies have been concerned with the first pairing. For example, Pavlov’s famous dogs received a sound stimulus (E1) and meat powder (E2). At first the dogs salivated only when meat powder (E2) was presented, but they soon started salivating at the sound of the bell (E1).

Trial and error learning, on the otherhand, is captured by the third pairing where the action of pressing a lever leads to the event of food being delivered (E2).

3.1.1 Learning E1 → E2 and E1 → no E2 associations

There is no doubt that animals can learn the first two associations. For example, Wasserman et al. (1974) put hungry pigeons in an operant chamber with a small disc on the right and left sides. E1 consisted of illumination of either disc with a white light for 10 seconds E2 was the presentation of food mid-way between the two discs.

Group 1 were exposed to E1 → E2 because food was presented immediately the disc light went off. Left and right discs were lit equally often and at irregular intervals.

Group 2 were exposed to E1 → no E2 by presenting food and lights as frequently but never together.

The behaviour measured was whether the pigeon tended to approach or withdraw from a disc when it was lit. The chamber was divided into two halves and the proportion of time the bird spent on the same side as the light was measured. Thus, a score of 0.5 indicates that a bird spent half the time on the side with the light. More than 0.5 shows a tendency to

approach the light, less than 0.5 shows a tendency to avoid it.

Figure 1 shows the approach-withdrawal scores. There was strong learning to approach when $E1 \rightarrow E2$ and weaker learning to avoid the light when $E1 \rightarrow \text{no } E2$.

Note: we are not claiming that pigeons learned that the light caused food, or that the light caused no food. What the animals actually learned is discussed later.

Nor are we concerned with why they learn.

For present purposes it is enough to say that they clearly learn something about the relationship between light and food.

As Fig. 1 shows, there was a third group. This had the light only and stayed at 0.5.

3.2 The problem of behavioural silence

Although we can be sure a behavioural change shows that learning has occurred we cannot assume that nothing is learned in the absence of a behavioural change.

If $E2$ is significant in motivation of species, then $E1 \rightarrow E2$ is hardly ever silent.

But $E1 \rightarrow \text{no } E2$ can often be silent. For $E1 \rightarrow \text{no } E2$ to affect behaviour, the animal must expect food to be given in the context of the experiment. It will not detect ‘no $E2$ ’ unless it expects $E2$ from the context. Look at Fig. 1 again. Both the Light-alone treatment and the Light + No food were examples of ‘no food’ but the pigeons only learned that the light meant no food when food was on the agenda.

3.2.1 Conditioned suppression

There is a technique, ‘Conditioned suppression’, which can be used to show silent learning in frightening $E1 \rightarrow E2$ situations (Rescorla, 1969).

First a rat is exposed to an $E1 \rightarrow E2$ association where $E1$ is a neutral stimulus, such as a tone, and $E2$ is an aversive stimulus such as a shock. The tone soon puts the animal in a state of fear. The fear induced by the tone puts the rat off its appetite so, if the tone is presented while the animal is pressing

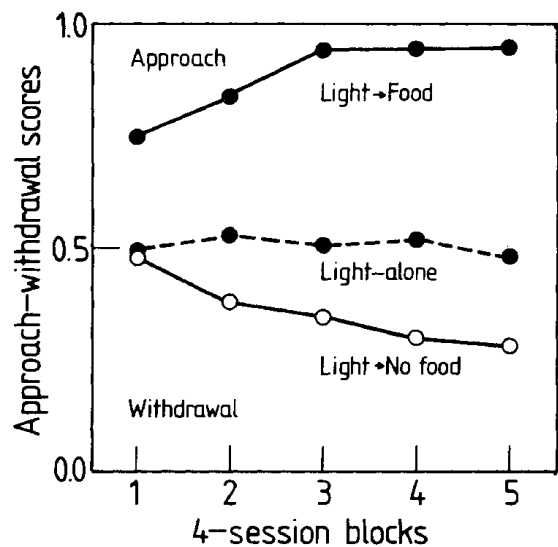


Figure 1: The acquisition of a tendency either to approach or to withdraw from a lighted key when pigeons are exposed to different correlations between illumination of the key and delivery of food. An approach-withdrawal score greater than 0.5 indicates a tendency to approach the illuminated key and less than 0.5 a tendency to withdraw. In the light \rightarrow food condition, illumination of the key was paired with food; in the light \rightarrow no food condition, the food never occurred during illumination of the key, nor shortly after it and in the light-alone condition the pigeons were simply exposed to illuminations of the key but were never presented with food. (After Wasserman *et al.*, 1974).

a lever to get food, or drinking from a water spout, the rate at which the rat presses or licks is slowed down. The *degree of suppression is a useful measure of how much the animal has learned about the tone shock relationship.*

The suppressive effects of a stimulus are usually measured by a ‘suppression ratio’ which compares the rate of responding during the stimulus with the rate immediately before. The suppression ratio is usually calculated as:

$$\frac{A}{A+B} \quad (1)$$

where A is the rate during the stimulus and B the rate just before the stimulus is presented. This ratio is 0.5 if the stimulus has no effect. If the stimulus produces suppression the ratio drops below 0.5 until it reaches a value of zero with complete suppression. So a score of zero means the fear induced by the tone is enough to stop feeding or drinking completely.

3.2.2 Summation test

Another way to reveal behaviourally silent learning is by a summation test. Suppose that after learning about a tone \rightarrow no shock association, the animals had learned about a light \rightarrow shock relationship.

When we now present the light alone it should suppress responding, but what would we expect if light and tone are presented simultaneously?

Under these circumstances the information about the tone \rightarrow no shock relationship is relevant because the animal expects the shock when the light comes on. Presentation of the tone should counteract the effect of the light. Rescorla demonstrated that this is just what happens

In conclusion, animals can learn both E1 \rightarrow E2 and E1 \rightarrow no E2. When E2 is motivationally significant then E1 \rightarrow E2 leads to changes in behaviour. By contrast, E1 \rightarrow no E2 can often be behaviourally silent and fail to affect performance directly.

The fact that something has been learned can still be revealed by:

1. Measuring the *suppression ratio*

2. Using a *summation test* to see if presentation of E1 will counteract the effect of another event predicting E2

3.3 Three questions about learning

Animal learning is selective and the first step in understanding the process is to specify the conditions for learning. Dickinson suggested that it is set up to detect and store information about causal relationships. If so, learning is likely to occur in conditions where a causal relationship is likely. The next section discusses whether these conditions are necessary for learning to occur.

A second question is about the way knowledge is represented. Is knowledge ‘declarative’, or ‘procedural’: giving someone directions to turn left at the pub, right at the phone box, and so on, is an example of procedural knowledge. Giving them a map is an example of declarative knowledge.

Declarative knowledge is easier to integrate, and it is easier to vary behaviour (or your route) on the basis of declarative knowledge. One set of roadworks which prevents you following procedural directions and you are lost.

Declarative knowledge is more flexible and robust but it cannot be observed unless it is being used. So we will have to be clever to demonstrate that animals store and use declarative knowledge. That is the subject of a later section.

Finally, we are interested in the mechanisms of learning. The animal selects information to be stored. How is the selectivity of the learning process achieved?

Additional questions include: How is the information stored? How is it recalled? These are discussed briefly by [Manning and Dawkins \(1992\)](#).

But in these two lectures I shall concentrate on how new knowledge is acquired.

4 Conditions of learning

If learning has evolved to enable animals to detect and store information about real causal relationships in their environment, then the conditions under

which learning occurs should be those where E1 is likely to be the cause of E2.

Figure 1 shows learning during pairings of E1 and E2. But *pairing is not enough to signal that E1 is a cause of E2*. Many irrelevant stimuli are paired with E2 by chance. So animals have to use a variety of evidence in deciding whether a particular pairing is a real relationship. We shall discuss three sources of information:

1. Event correlation.
2. Causal relevance of pairing: ('feeling sick' is more easily paired with taste stimuli than with sounds).
3. Temporal relationship.

4.1 Event correlation

If E1 is the cause of E2, *then E2 should not occur without E1*. This is true of the pattern shown in Fig. 2A. But Fig. 2B shows pairings of E2 when E1 as well as pairings of E2 when no E1. In this case, pairings should be put down to chance even though E1 \rightarrow E2 pairings occur with exactly the same frequency as in Fig. 2A.

So, the amount an animal learns about E1 \rightarrow E2 should be influenced by the number of times that E2 occurs without E1.

Rescorla (1968) was the first to test this idea systematically. He used a suppression procedure.

E1 was a tone, E2 a shock. Rats were first trained to press a lever in an operant chamber, then transferred to a conditioning chamber. Here rats received two-minute tones.

There were four groups as well as the controls. Each group received a shock in four out of ten tone presentations so the probability of a shock, given the tone was 0.4: $P(\text{shock when tone}) = 0.4$. In otherwords, all groups (except the controls) had the same number of pairings of E1 \rightarrow E2, just as in Fig. 2.

Rescorla then varied the probability of a shock occurring between tones. In otherwords he varied $P(\text{shock when no tone})$. The results are shown in Fig. 3.




Figure 2: The pattern of events presented when there are different correlations between E1 and E2. In Panel A, E2 only occurs during or shortly after E1 and the two events are positively correlated. In Panel B, E2 is just as likely to occur when E1 is absent as when it is present and the two events are uncorrelated. So in Panel A, $P(E2 \text{ when } E1)$ is high and $P(E2 \text{ when no } E1)$ is low. In Panel B, $P(E2 \text{ when } E1) = P(E2 \text{ when no } E1)$.

1. The control group never received a shock and they showed a suppression ratio of 0.5 which means zero suppression.
2. Group '0.4' had $P(\text{shock when no tone}) = 0.4$, so there should be no evidence of a causal relationship since the shock is just as likely to occur in presence or absence of a tone. They showed no suppression, like the controls.
3. Group '0.2' had $P(\text{shock when no tone}) = 0.2$, so the shock was twice as likely to occur after a tone as after no tone. They showed some suppression.
4. Group '0.1' had $P(\text{shock when no tone}) = 0.1$, so the shock was four times as likely to occur after a tone as after no tone. They showed very significant suppression.
5. Group '0.0' had $P(\text{shock when no tone}) = 0.0$, so the shock always occurs after a tone and never without one. The suppression was almost complete.

The results in Fig. 3 were obtained when the rats were put back into the operant chamber after 60 tone presentations.




Figure 3: The degree to which a tone suppressed lever-pressing for food after various groups of rats had received different correlations between the tone and a shock. The different correlations were brought about by holding the $P(\text{shock when tone})$ constant at 0.4 for all groups and varying $P(\text{shock when no tone})$ across the different groups (After Rescorla, 1968). Note the suppression ratio is $A/(A + B)$ where A is the rate of lever presses during the tone and B is the rate just before the tone

If learning depended only on pairing, then all except the control group should have shown the same level of suppression, since they all received the same number of pairings.

CONCLUSION: For an animal to learn a tone \rightarrow shock relationship, then $P(\text{shock when tone})$ must be greater than $P(\text{shock when no tone})$ and the amount the animal learns depends on this discrepancy. What Rescorla has shown is that varying $P(\text{shock when no tone})$ varies the intensity of learning and when we come to think about mechanisms this may be a problem. It is easy to imagine how $P(\text{shock when tone})$ is detected and learned by the nervous system, because they both occur at the same time, but *how is the value of $P(\text{shock when no tone})$ measured?* That is the key to understanding how animals detect a correlation between events.

4.2 Overshadowing

Up to now it looks as if the shock might have only one cause but, from the animal's point of view, there are a myriad of potential causes in any environment. There are other stimuli provided by the conditioning chamber, as well as those due to the animal's own behaviour. Does the rat blame itself for the shocks?

The rat has to learn which of the events and cues is the most likely cause or predictor.

Unfortunately, we cannot eliminate contextual cues, but we may be able to create and control a second stimulus.

Mackintosh (1976) trained a rat to press a lever for food, and then presented a series of learning trials while the animals were pressing it. Group 'L' received a light stimulus while the rest received light plus noise. Low noise was 50 dB, high noise was 85 dB. A shock was presented at the end of the stimulus on all trials.

After this training, the light alone was presented during bar pressing to measure the suppression ratio.

Figure 4A shows that feeding of the 'L' group, which had experienced light alone, was almost completely suppressed by the light stimulus. There was very much less suppression in the 'LN' group. This

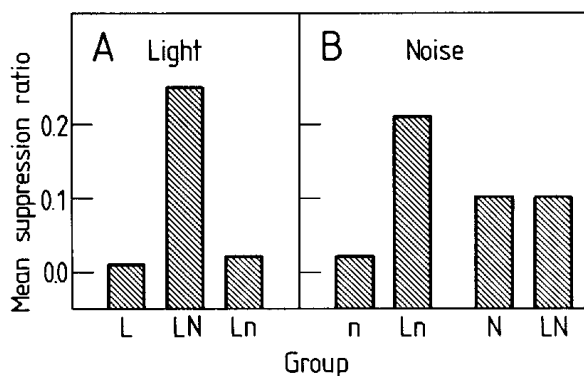


Figure 4: The degree to which a light (Panel A) and a noise (Panel B) suppressed lever-pressing for food after different groups of rats had received pairing of a shock and the light (L), a weak noise (n) and an intense noise (N), a compound of the light and a weak noise (Ln), or a compound of the light and intense noise (LN). (After Mackintosh, 1976.)

is consistent with the idea that the rats had learned to blame the shock on the noise when the noise was very intense. *This is an example of overshadowing*: the loud noise overshadowed the light when both were presented together during training.

The ‘Ln’ group did not seem to blame the shock on the low noise stimulus and their feeding was almost completely suppressed by the light, so the low noise did not overshadow the light appreciably.

CONCLUSION: the noise overshadowed the light and the degree of overshadowing depended on the intensity of the noise.

Mackintosh also asked whether the presence of the light (during training) influenced the effect of the noise. He presented the noise alone (Fig. 4B) after training and found that the presence of the light *during training* overshadowed the low noise (the ‘Ln’ group was less suppressed by noise than the ‘n’ group). But it did not overshadow the loud noise (‘N’ and ‘LN’ groups were equally suppressed).

So this is evidence that two or more EIs, presented together, compete for the learning mechanism. The relative intensity of the two stimuli will determine which is learned most effectively and how much one

will overshadow the other.

Since all stimuli compete with background cues, it may be that learning is always less than some theoretical maximum. In other words, all learning about a given stimulus is reduced to some extent because background stimuli compete with it for the learning mechanism.

4.3 Overshadowing by background cues

I mentioned above that, when it comes to thinking about mechanisms, it may be hard for the animal to measure $P(\text{shock no tone})$ and the idea of overshadowing may help. You need to take the next bit slowly, and may be several times.

In Rescorla (1968) the uncorrelated group ‘0.4’ received more shocks than the correlated group ‘0.0’. (This is as in Fig. 2A.) Now background stimuli were always present, so it may be that the background stimuli came to be associated with the shock more than did the tone, and then these background stimuli came to overshadow the tone.

Putting this another way, the top two rows of Table 1 show the stimuli that were present in Rescorla’s experiment. *A* stands for the target stimulus, which was a tone in Rescorla’s experiment. *B* stands for the background or contextual cues. When the tone was presented, the rats experienced *A* and *B* together but, in between successive tone presentations, they received *B* alone. Animals in both groups received the same number of pairings of *A* + *B* and shock. The only difference was that ‘*B* alone’ was paired with the shock in the uncorrelated group (Table 1). *So the rats in the uncorrelated group were right to blame the shock on B, and not on A. B was always present when the shock occurred so perhaps B came to overshadow A.*

We cannot test this directly, but we can have two controllable stimuli, *A* and *B*, and run an experiment exactly as Table 1 sets out. Remember, the top two rows of Table 1 were a summary of Rescorla’s experiment in which the background was uncontrolled. But the table can be used to design an experiment in which one of the background stimuli is controlled and presented like *B* in the table.

Table 1: The pairings which occur in different E1–E2 correlations

Condition	Type of E1 presentation	
	<i>A</i> and <i>B</i>	<i>B</i> alone
Correlated	E2	no E2
Uncorrelated	E2	E2
Simple overshadowing	E2	

Saavedra (reported by [Wagner, 1969](#)) did just such an experiment, using a rabbit eyeblink reaction. He used a short 1-second tone or light as E1 and a mild shock to the eye as E2. The shock automatically elicits an eyeblink and, when the rabbit learns about the $E1 \rightarrow$ shock association, then E1 elicits a blink. Saavedra used the layout in Table 1 using a light as stimulus *A* and a tone as *B*. All animals received the same number of $AB \rightarrow$ shock pairings.

After this training, Saavedra presented test trials with AB together, *A* alone or *B* alone. Figure 5 shows the percentage of tests eliciting an eyeblink. The higher the percentage, the more the animals had learned about the association being tested. The interesting result is that the rabbits in the correlated group learned more about the $A \rightarrow$ shock association than in either of the other two groups (compare the ‘A’s in Fig. 5).

So the $B \rightarrow$ shock pairings in the uncorrelated group helped *B* to overshadow *A*. The uncorrelated group did, indeed, learn to blame the shock on *B* and so were less likely to blame it on *A*.

Similar results have been found in other experiments and species and, in general, it seems that in the correlated condition the animals blame the shock on the tone and not on background cues. The reverse is true in the uncorrelated condition and the animals blame the shock mainly on the background.

4.4 Action \rightarrow E2 learning

It is impossible to do the same experiments when E1 is an action because the frequency of E1 is not under the experimenters control. The whole




Figure 5: The percentage of test trials on which the presentation of stimulus A, stimulus B, and a compound of stimuli A and B elicited an eyeblink after different groups of rabbits had been exposed to the uncorrelated, simple overshadowing or correlated training schedules (see Table 1). (Saavedra reported by Wagner, 1969)

thrust of these experiments is that the same number of $E1 \rightarrow$ E2 pairings occurs across treatments. Several attempts have been made to overcome this problem and the results are similar (eg: Fig. 6) ([Mackintosh and Dickinson, 1979](#)).

4.5 E1 \rightarrow no E2 learning

Typically, this kind of learning is silent and has to be exposed by careful design but there is no doubt it occurs ([Dickinson, 1980](#), pp. 39–42). The key difference is that $P(\text{shock when no E1})$ has to be greater than $P(\text{shock when E1})$ which is the opposite of the conditions for $E1 \rightarrow$ E2.

4.6 Overview of conditions for learning

Animals appear to learn an $E1 \rightarrow$ E2 relationship if $P(E2 \text{ when } E1)$ is greater than $P(E2 \text{ when no } E1)$. The opposite is necessary to learn an $E1 \rightarrow$ no E2 relationship.

Figure 7 shows a space representing every possible combination of $P(E2 \text{ when } E1)$ and $P(E2 \text{ when no } E1)$. The uncorrelated condition is on the diagonal




Figure 6: The rate at which different groups of rats ran in a wheel after having been exposed to the uncorrelated (U), simple overshadowing (O), or correlated training schedules (C) with running as E1. (After Mackintosh and Dickinson-1979.)




Figure 7: A contingency square representing possible combinations of the conditional probabilities of E2 occurring when E1 is present, $P(E2/E1)$, or absent, $P(E2/noE1)$. Along the diagonal E1 and E2 are uncorrelated and associative learning does not appear to occur. Above the diagonal E1 and E2 are positively correlated and $E1 \rightarrow E2$ learning occurs, whereas below the diagonal the two events are negatively correlated and $E1 \rightarrow noE2$ learning is observed. The filled circles represent the conditions run in Rescorla's (1968) experiment (see Fig. 3).

and no learning occurs if the two probabilities are equal. $E1 \rightarrow E2$ learning occurs above the diagonal and $E1 \rightarrow noE2$ learning occurs below it. The four filled circles show the probabilities used by Rescorla (1968) and learning did indeed increase with distance from the diagonal.

It is important to be clear that our measures of the intensity of learning are indirect and are not related in a simple linear way to the probabilities in Fig. 7. We have no idea of the shape of the lines that connect point of equal learning throughout this space.

4.7 Causal relevance

In the real world, causal chains are detected partly by correlation between events and partly by the type

of event. If you feel suddenly sick after visiting a new restaurant you are likely to blame the food not the wallpaper or music, and yet wallpaper and music were present throughout the meal.

If learning mechanisms are designed to detect and store information about causal relationships, an animal should learn more readily about the taste \rightarrow illness association than the relationship between say a tone or light and illness.

Domjan and Wilson (1972) gave two groups of rats a 35-second oral infusion of saccharin. Immediately after this infusion, one group was injected with lithium chloride to induce sickness and the other with saline as a control. Another pair of groups was given a 35 second exposure to a buzzer and then injected with lithium chloride or saline.

After three such pairings the rats were given a preference test with two water spouts. For the first pair of groups the choice was between pure water and saccharin solution. The group made sick avoided the saccharin and preferred the water. Those with the saline injections preferred the saccharin solution (Fig. 8).

The buzzer group were given the choice between a spout which activated a buzzer and one which did not. There was no significant difference between those made sick and those that received saline injections.

Thus, the rats were more able to learn an association between sickness and taste than between sickness and a buzzer.

However, this might be because saccharin is a stronger stimulus than the buzzer. To test this, Domjan and Wilson (1972) tested a third pair of groups with a mild shock and found that the buzzer was a stronger stimulus than saccharin in learning about shocks (compare saccharin \rightarrow shock with buzzer \rightarrow shock in Fig. 8). Thus, the type of stimulus is relevant in learning about different consequences and animal learning is selective, at least under some circumstances.

4.8 Temporal relationships between events

We have seen that there has to be a positive correlation between events for rats to learn about the




Figure 8: The preference of rats for drinking from a water spout when the animals were given a choice between this water spout and either one delivering a saccharin solution (Panel A) or another water spout which activated a buzzer when the animal made contact (Panel B). The various groups received prior training in which either the taste of the saccharin solution (Panel A) or the buzzer (Panel B) was paired with the induction of illness by an injection of lithium chloride, a shock, or a simple saline injection. (After Domjan and Wilson, 1972.)




Figure 9: The suppression of licking produced by a tone after different groups of rats had experienced a single pairing of a tone and shock with various intervals between the onset of the tone and the shock. The control group just received a presentation of the shock but no tone. The degree of suppression was measured by recording latency to complete 10 licks in the presence of the tone on test; the longer the latency the more severe the suppression. (After Mahoney and Ayres, 1976.)

association. In doing this we have assumed that the events have to be close together in time. In addition, in all cases discussed so far, the onset of E1 has occurred before E2. Are these timings necessary?

There have been many experiments in which timing of the events has been varied. For example, Mahoney and Ayres (1976) gave rats a single presentation of a 4-second tone and a 4-second shock and varied the interval between the onset of these two stimuli for different groups. For four groups, the onset of the tone preceded the shock by 0, 4, 8 or 150 seconds. For two groups it lagged the shock by 4 or 8 seconds. The next day they used a suppression test, measuring the time to complete ten licks of a spout after the onset of the tone. The stronger the learning the longer ten licks should take (latency increased).

As Fig. 9 shows the strongest learning was found when the shock was delivered 8 seconds after the tone. Presentations before the tone were not learnt.

It is a cardinal feature of our idea of cause and effect that the cause must come first, but it is important to remember that the rat is using events as

predictors or indicators. It is not able to identify true causes.

One nuance of this concept is that the amount the rat learns about the background cues should vary inversely with the amount it learns about the actual tone. Odling-Smee (1975) analysed this part of rat behaviour. Their conclusion was that as time increases between the two events there are more and more other potential cues that the rat may blame the shock on, so the effectiveness of the tone becomes masked by background cues. Again, the mechanism is consistent with the idea that potential cues compete for the learning mechanism.

If it is correct to suppose that variations in time intervals between two events affect learning through the overshadowing process we should expect such variations to have far less effect in the case of taste \rightarrow illness learning than in the case of tone \rightarrow shock learning. Smith and Roll (1967) found appreciable learning when the interval between stimulus and illness was up to 12 hours.

Revusky (1971) interposed other tastes between the saccharin and the illness and found that the rats blamed the illness on the later stimulus, depending also on its concentration.

Dickinson argues that it is reasonable that rats associate tastes and illness over long delays because the environment provides few intervening background stimuli which are relevant. By contrast, for tone \rightarrow shock learning the environment contains a host of potentially relevant stimuli so there are many more background cues on which to blame the shock. This is no doubt true, but there may be other processes in the nervous system that cause a degradation of ‘temporary memory’ with time and the long intervals shown in taste \rightarrow illness learning may use different types of system.

At the very least, the analysis has revealed alternative theories that may be compared in future research.

4.8.1 Temporal events in honey bee learning

A remarkable example of sensitivity to time between events is provided by the honey bee (Menzel and Erber, 1978). Honey bees learn and re-




Figure 10: Strong association between colour and reward is dependent on exact timing. As the bee approaches the disc its colour has no significance until about 4 s before the bee begins to suck sugar. Then there is a sharp peak of strong association which has disappeared within 2 s following this onset of sucking. Apart from this short time, the colour the bee sees during the approach or for the whole time it takes to fill its crop, is of no importance. (From Menzel and Erber, 1978)

tain the effects of a single association between flower colour and food reward. Since they will visit and feed from a dish of sugar solution at a table it is possible to illuminate the dish from below and change its colour at will. It takes a bee a minute or two to feed and the colour of the ‘flower’ can be changed during this interval. Menzel and Erber (1978) changed flower colour between yellow and blue at various stages and then gave the bees a choice to see which colour they had learned to link with sugar. The results showed that the period from 4 sec before to 2 seconds after the tongue first touched the sugar solution was the most sensitive (Fig. 10). If the flower was initially yellow but changed to blue for this brief period the honey bee would subsequently prefer blue flowers.

It is curious that this window of sensitivity is similar in length to that shown by rats, although the mechanism may be different.

4.9 Summary

In Section 4 we have pointed out that the animal has to select from a host of competing stimuli those

which are the best predictors of significant effects. The problems are the same whether these are the causes, or simply indicators.

Simple pairing of events is not enough because some events will be paired by chance. Learning that a tone predicts a shock should and does occur when $P(\text{shock when tone})$ is greater than $P(\text{Shock when no tone})$.

However, this raises problems in understanding the mechanisms because it is not easy to see how $P(\text{shock when no tone})$ is measured. It turns out that rats tend to blame the background cues when this probability is high because when $P(\text{shock no tone})$ is high there are many pairings of background and shock. The background cues then overshadow the tone which blocks learning of the tone \rightarrow shock association when it would not be appropriate.

Hence, learning of correlations can be understood in terms of the two types of pairing: pairing of the tone and shock and pairing of background cues and shock. Whichever is more frequent tends to win.

Overshadowing by background cues may also explain the importance of close timing between events when an association is learned. The longer the gap between tone and shock, the more chance background cues have of intervening and hence blocking the tone shock learning. It may, instead, be that a direct effect of delay occurs, and the jury is out on that.

The causal relationship between events is clear in some cases: a rat is more likely to associate illness with a taste than with a buzzer, but it is more likely to associate a shock with a buzzer than with a taste. No doubt such selectivity is genetically controlled and the product of natural selection.

These common principles appear to apply to learning of $E1 \rightarrow E2$ and $E1 \rightarrow \text{no } E2$, whether $E1$ is an external stimulus or the animal’s own actions. Association of $E1 \rightarrow E2$ is almost always visible if $E2$ is a significant effect such as a shock or food. Learning of $E1 \rightarrow \text{no } E2$ is commonly silent and has to be demonstrated by special techniques such as suppression or summation tests.

Table 2: Integration of associative representations

Association	Procedural representation	Declarative representation
tone \rightarrow food	‘When tone is on approach magazine’	‘the tone causes food’
food \rightarrow illness	‘When food is present suppress eating’	‘the food causes illness’
integration		‘the tone causes illness’

Table 3: Design of Holland and Straub (1979) experiments

Group	Stage 1	Stage 2	Test
E	noise \rightarrow food	food \rightarrow LiCl	noise
C	noise \rightarrow food	LiCl	noise

5 Associative representations

5.1 Declarative and procedural representations

The difference between declarative and procedural knowledge has been outlined above and further examples are given in Table 2. Declarative knowledge is easier to integrate but it is usually silent until it is used. Experiments therefore have to be designed to prove the existence of declarative knowledge.

Holland and Straub (1979) designed an experiment based on Table 2. The design is shown in Table 3. All rats received a series of sessions in which a 10 second noise stimulus preceded delivery of two food pellets. Figure 11 shows that there was an immediate rise in approaching the food magazine when the tone was presented.

In the second stage, the rats in Group E were given 50 pellets in their home cage and then injected with lithium chloride to make them ill. This was repeated four times on different days. By the fourth day there was a dramatic fall in the number of pellets the rats consumed, showing they had learned about the food

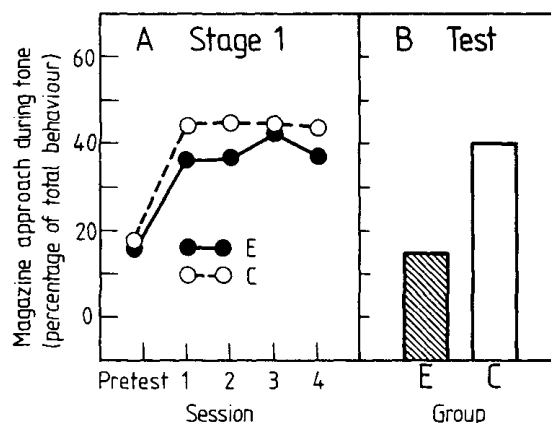


Figure 11: Panel A: The acquisition of magazine approach during a tone which was paired with the delivery of food to the magazine. The strength of magazine approach is expressed by the percentage of all behaviours observed during the tone which consisted of approaching the magazine. Panel B: The strength of magazine approach after the food had been paired with the induction of illness in Group E but not Group C (see Table 3). (After Holland and Straub, 1979.)

illness association. The control group simply received injections of lithium chloride with no food. This is an example of a ‘devaluation’ procedure and, in practice, [Holland and Straub \(1979\)](#) used a more sophisticated devaluation procedure but the details are not necessary here.

In the final test, the rats were returned to the experimental chamber and received a number of presentations of noise. If the rats in Group E were able to integrate the information then they should approach the magazine less than at the end of stage 1. [Figure 11](#) shows that they did, while the control group continued to approach the magazine at the original rate. Clearly, the declarative model fits best in this case.

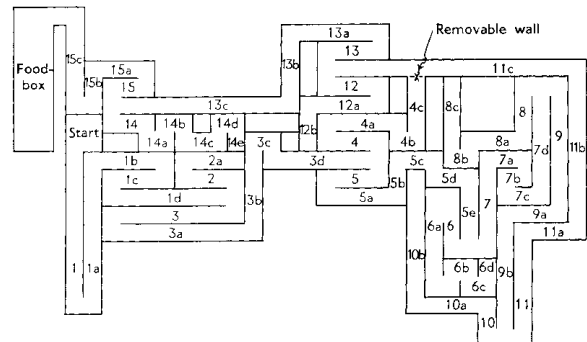
Nevertheless, integration does not always occur. [Holland and Rescorla \(1975\)](#) extended the chain of integration to three steps and rats succeeded in integrating the information. Other integration, however, failed. Taken at face value their results suggest that the rats’ ability may be limited, but interpretation is difficult because it depends on the kind of task that the rats were being asked to do. At present we do not really know why successful integration occurs in some cases but not others.

It seems almost certain that both types of representation occur and the role of each depends on how automatic a task has become. When learning to drive the operations of changing gear may be declaratively stored at first but, with experience the procedures become so familiar that a procedural representation takes over. Similar, procedural, habit formation no doubt occurs in many animal species.

5.2 Declarative knowledge in maze behaviour

[Manning and Dawkins \(1992\)](#) give a dramatic example of geographical knowledge in rats ([Fig. 12](#)).

In sharp contrast to the rats, [Lorenz](#) describes water shrews that learned the geography of their aquarium, jumping from island to island in the dark. They relied entirely on proprioception and repeated the virtually the same steps and jumps for each journey. So



[Figure 12](#): A complex maze used by [Shephard](#) to test ‘reasoning’. After rats have learned the maze, the section indicated by X is removed, changing a blind alley into a short cut. Having discovered the change whilst running along 11c, and exploring from there a little into 4c, some rats entered 4 (and thence 4a, 4b, 4c) instead of 5 on the next trial (from [Maier and Schneirla 1935](#)).

fixed was their behaviour that they would leap towards islands that they could see had been removed. Bats, too, navigate their caves by procedural knowledge and do not bother to echo-locate once they know the route. This allows them to be caught by putting up wooden boards and the bats fly straight into them.

Hence, there is clear evidence of both procedural and declarative forms of knowledge, used by different species in ways which are usually appropriate.

5.3 A very simple model of learning

One feature of learning an $E1 \rightarrow E2$ is particularly striking: once learned, the response to $E1$ appears to be the same as the normal response to $E2$. It therefore appears that there has been a stimulus substitution upstream of the system controlling the response. For example, [Jenkins and Moore \(1973\)](#) noted the difference between the way pigeons responded to food and water. When eating grain pigeons make short, sharp, pecking movements in which the beak is open at the point of contact and the eyelids almost closed. When drinking, by contrast, the bird places its beak in the water, slightly open, and swallows with its eyes fully open. When trained in an operant chamber the




Figure 13: A schematic illustration of the possible excitatory-link formed after exposure to $E1 \rightarrow E2$. The S-S (Stimulus-Stimulus) link is declarative ($E1$ comes to predict $E2$) while the S-R (Stimulus-Response) link is a procedural representation ($E1$ triggers pecking directly). The observed behaviour is an amalgam of the influences of the aroused $E1$ and the response to $E2$.

pigeons responded to keys in the same way appearing to ‘drink’ the water key and ‘eat’ the grain key.

Figure 13 provides a framework for discussion and shows a schematic representation of a declarative model of stimulus substitution, originally proposed by Pavlov (1927) and developed by Konorski (1948); Dickinson and Boakes (1979). Dickinson makes much of this model and of stimulus substitution, but I am not sure how far they help. The results described so far seem rigorous, while the discussion of the model and of stimulus substitution seem rather weak.

The direct links of this model should not be thought of as connections formed in the nervous system, and it is clear that it could not handle the complex maze behaviour shown in Fig. 12. Indeed, the processes discussed here are still a long way from mechanistic explanation, but the demonstration of silent learning has been a great step forward in setting the agenda for physiological explanations of the learning mechanisms.

6 Mechanisms of learning

The chapter on mechanisms of learning contains a great deal of rigorous argument and data but it is beyond the scope of two lectures. Dickinson discusses

two main theories and I give a brief account of the Rescorla-Wagner theory to demonstrate its simplicity.

The central idea of the theory is that the degree to which $E1$ and $E2$ become associated depends on the extent to which they are processed together. The change in learning at each presentation is given by dV and

$$dV = (\text{processing of } E1)(\text{processing of } E2) \quad (2)$$

In the original version of the theory the processing of $E1$ is seen as fixed and does not change with learning. That of $E2$ depends on how surprising it is. As learning occurs, $E2$ becomes predicted more strongly and is therefore less surprising, so dV becomes smaller with each presentation.

This equation can be rewritten as

$$dV = \alpha\beta(\lambda - V) \quad (3)$$

where

α is intensity of $E1$

β is intensity of $E2$

λ is initial value of processing of $E2$ (when it is at its most surprising)

As learning occurs, V approaches λ until no further learning occurs.

This equation can be extended in a simple way to handle background stimuli and the theory explains most of the experiments described in the book. In particular, it explains overshadowing and blocking of one association by another. However, it fails to explain variations in the processing of $E1$ and Wagner later extended the theory to include processing of $E1$. There remain one or two phenomena that the theory does not explain adequately and more work is needed.

A number of authors feel that the central idea of the Rescorla-Wagner theory is unintuitive: that learning depends on how surprising $E2$ is in the context of $E1$. They feel, instead, that learning of $E1 \rightarrow E2$ should depend on how good a predictor $E1$ is of $E2$, rather than how surprising $E2$ is. A group of theories have developed around these ideas, and are

well reviewed by Dickinson, but they are less easy to describe and less well developed than the Rescorla-Wagner theory. Time and further experiments will help to choose between them or, more likely, develop stronger alternatives.

A more upto date review, which follows on from these notes is provided by [Pearce \(1997, Chapter 3\)](#).

References

- A Dickinson. *Contemporary animal learning theory*. Cambridge University Press, Cambridge, 1980.
- A. Dickinson and R.A. Boakes, editors. *Mechanisms of learning and motivation: a memorial volume to Jerzy Konorski*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1979.
- M. Domjan and N.E. Wilson. Specificity of cues to consequence in aversion learning in the rat. *Psychonomic Science*, 26:143–145, 1972.
- P.C. Holland and R.A. Rescorla. The effect of two ways of devaluing the unconditioned stimulus after first- and second-order appetitive conditioning. *Journal of Experimental Psychology: Animal Behaviour Processes*, 1:355–363, 1975.
- P.C. Holland and J.J. Straub. Differential effects of two ways of devaluing the unconditioned stimulus after Pavlovian appetitive conditioning. *Journal of Experimental Psychology: Animal Behaviour Processes*, 5:65–78, 1979.
- H.M. Jenkins and B.R. Moore. the form of the auto-shaped response with food or water reinforcers. *Journal of the Experimental Analysis of Behaviour*, 20:163–181, 1973.
- J. Konorski. *Conditioned reflex and neuron organisation*. Cambridge University Press, Cambridge, 1948.
- N.J. Mackintosh. Overshadowing and stimulus intensity. *Animal Learning and Behaviour*, 4:186–192, 1976.
- N.J. Mackintosh and A. Dickinson. Instrumental (Type II conditioning). In A. Dickinson and R.A. Boakes, editors, *Mechanisms of Learning and Motivation*, pages 143–170. Lawrence Erlbaum Associates, Hillsdale, N.J., 1979.
- W.J. Mahoney and J.J.B. Ayres. One-trial simultaneous and backward fear conditioning as reflected in conditioned suppression of licking in rats. *Animal Learning and Behaviour*, 4:357–362, 1976.
- N.R.F. Maier and T.C. Schneirla. *Principles of animal psychology*. McGraw-Hill, New York, 1935.
- A. Manning and M.S. Dawkins. *An introduction to animal behaviour*. Cambridge University Press, Cambridge, 4th edition, 1992.
- R. Menzel and J. Erber. Learning and memory in bees. *Scientific American*, 23 (1):80–87, 1978.
- F.J. Odling-Smee. Background stimuli and the inter-stimulus interval during Pavlovian conditioning. *Quarterly Journal of Experimental Psychology*, pages 387–392, 1975.
- I.P. Pavlov. *Conditioned reflexes*. Oxford University Press, Oxford, 1927.
- J.M. Pearce. *Animal learning and cognition*. Psychology Press, 1997.
- R.A. Rescorla. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Conditioning*, 66:1–5, 1968.
- R.A. Rescorla. Conditioned inhibition of fear resulting from negative CS–US contingencies. *Journal of Comparative and Physiological Psychology*, 67:504–509, 1969.
- S.H. Revusky. The role of interference in association over delay. In W.K. Honig and P.H.R. James, editors, *Animal Memory*, pages 155–213. Academic Press, New York, 1971.
- R.C. Rizley and R.A. Rescorla. Associations in second-order conditioning and sensory preconditioning. *Journal of Comparative and Physiological Psychology*, 91:1–11, 1972.

J.C. Smith and D.L. Roll. Trace conditioning with X-rays as the aversive stimulus. *Psychonomic Science*, 9:11–12, 1967.

W.H. Thorpe. *Learning and instinct in animals*. Methuen, London, 1963.

A.R. Wagner. Stimulus selection and ‘modified continuity theory’. In *The psychology of learning and motivation*, volume 3, pages 1–43. Academic Press, New York, 1969.

E.A. Wasserman, S.R. Franklin, and E. Hearst. Pavlovian appetitive contingencies and approach versus withdrawal to conditioned stimuli in pigeons. *Journal of Comparative and Physiological Psychology*, 86:616–627, 1974.