Pavlov and Associationism

Nicholas J. Mackintosh
University of Cambridge

Pavlov's contribution to experimental psychology was to invent a technique that allowed him to undertake a prolonged and systematic series of well-controlled experiments that, astonishingly enough, uncovered many if not most of the phenomena of what is rightly called Pavlovian conditioning. It was not for another 30 years or more that English-speaking psychologists began to match that achievement. Of course there have been new developments and discoveries since his time. Two examples are discussed: the important role of variable associability or attention even in simple conditioning, and the rigorous application of associative learning theory to the behavior of adult humans.
Keywords: Pavlov, associationism, attention, associative learning

La contribución de Pavlov a la psicología experimental fue inventar una técnica que le permitió acometer una serie prolongada y sistemática de experimentos bien controlados que desvelaron muchos, si no la mayoría, de los fenómenos de lo que justamente se llama condicionamiento pavloviano. Los psicólogos angloparlantes tardarían 30 años o más en igualar ese logro. Por supuesto, se han dado nuevos desarrollos y descubrimientos desde su época. Se comentan dos ejemplos: el importante papel de la asociabilidad variable o la atención incluso en el condicionamiento simple y la aplicación rigurosa de la teoría del aprendizaje asociativo a la conducta de los seres humanos adultos.
Palabras clave: Pavlov, asociacionismo, atención, aprendizaje asociativo

Pavlov's Contribution

The history of associationism, as of all good, and some not so good, ideas, can be traced back to ancient Greek philosophy. But it was Pavlov's discovery of the conditioned reflex that made the scientific study of associationism or associative learning possible. Pavlov's work has not been without its critics: George Bernard Shaw (1944, p. 208) famously asserted that "the story of his dogs in his book is a crackle of blazing nonsense from beginning to end." But this was after H.G. Wells had announced that if he saw Pavlov and Shaw drowning, and had only one lifebuoy available, he would throw it to Pavlov and leave Shaw to drown. Others have shown that it is possible to find earlier descriptions of people or other animals behaving in ways that we would now describe in the terminology that Pavlov invented. A nice example is provided by Rosenzweig (1962), who cites Robert Whytt's *An Essay on the Vital and Other Involuntary Motions of Animals*, published in 1763: "The remembrance or *idea* of substances formerly applied … produces almost the same effect as if these substances were really present. Thus the sight, or even the recalled *idea* of grateful food, causes an uncommon flow of spittle into the mouth of a hungry person; and the seeing of a lemon can produce the same effect in many people."

But it was still Pavlov who discovered conditioned reflexes, in the sense that it was the procedures and techniques of his experiments that identified the conditioned reflex as a phenomenon worth studying, showed how it could be studied, and made the serious and scientific investigation of associative learning possible. We are all in his debt.

Pavlov invented a technique for the study of conditioning, and used that technique in a prolonged and systematic series of well-controlled experiments. One hundred years later, it is easy for us to criticize many of his experiments, and to point out that they rarely meet the standards of control and rigor that the editor of the *Journal of Experimental Psychology* requires today. A hundred years ago, however, his experimental program represented an astonishing achievement, and one that it took English-speaking psychologists *at least* 25 to 30 years to match. It is true that few if any people today are studying salivary conditioning in dogs, but *no one* has studied cats learning to escape from Thorndike's puzzle box for more than half a century. Indeed, we now know, thanks to the work of Moore and Stuttard (1979), that Guthrie and Horton (1946), the last people to do so, were merely eliciting, and then carefully recording, their cats' greeting gestures. Pavlov's experimental paradigm set the standard for what is required for the experimental analysis of associative learning, and virtually all the advances that have been made since 1965 in our understanding of associative learning, even if they have not relied on salivating dogs, have been based on procedures closely modeled on those he developed.

That paradigm also allowed Pavlov himself to make a remarkable number of discoveries about conditioning that have stood the test of time—not a claim that anyone would want to make about Thorndike's puzzle box experiments. The list includes: acquisition, extinction and spontaneous recovery; forward and backward, delay and trace conditioning; higher-order conditioning; generalization and discrimination; positive and negative patterning (the A+, B+, AB− discrimination); conditioned inhibition and the summation and retardation tests needed to detect it; generalization of inhibition; counterconditioning; reinstatement; external inhibition and disinhibition; overshadowing as a function of relative salience or relative validity; and transfer along a continuum. (It is probably wise to pass over his experiments on positive and negative induction: some find a parallel in more modern work on contrast, but some of his results remain puzzling and not well replicated).

Pavlov repeatedly and adamantly insisted that he was a physiologist, not a psychologist, and much of his theorizing was couched in quasi-physiological language. Most of it has, deservedly, been rejected, and it was left to Konorski (1948) to reinterpret Pavlov's ideas in more acceptable physiological terms. Konorski's analysis is more acceptable, of course, because, rather than appealing to spreading waves of excitation and inhibition across the cortex, it talked of excitatory and inhibitory connections. In other words, it was an associative analysis—and Pavlov's account of conditioning was equally associationist.

Although he might not want to have admitted it, Pavlov did indeed engage in psychological theorizing. He drew a sharp distinction between excitatory and inhibitory conditioning. He attributed excitatory conditioning to the ability of a conditioned stimulus (CS) paired with an unconditioned stimulus (US) to activate a representation of (or the center for) that US, and thereby activate components of the unconditioned response (UR) elicited by the US itself. This stimulus-substitution theory, which can be contrasted with an S-R theory of conditioning, may need substantial qualification, but remains a viable and, at least in part, a widely accepted, account of Pavlovian conditioning to this day. He argued that extinction and inhibitory conditioning do not simply represent the weakening of an excitatory association, but rather the strengthening of a process of inhibition, which opposes or counteracts the still existing excitatory association. Again, this is unquestionably a psychological theory—and again one that would be widely accepted today. Only some rather tortuous argument, and refusal to consider all Pavlov's evidence, allowed Skinner (1938) to dismiss the concept of inhibition.

Just as important, Pavlov was surely right to complain about the way in which the conditioned reflex was used by many American psychologists, when he wrote of Guthrie: "The psychologist takes conditioning as the principle of learning, and accepting this principle as not subject to further analysis, not requiring ultimate investigation, he endeavors to apply it to everything" (Pavlov, 1932, p. 91).

Ultimate investigation was what Pavlov aimed to achieve, and what was surely needed.

Developments since Pavlov

I have been talking of experiments undertaken, and theories formulated, more than three quarters of a century ago. One would hope that the study of conditioning and associative learning had made some progress since then. It is surely no discredit to Pavlov to say that it certainly has. On the contrary, one of the best indicators of Pavlov's achievement as a scientific pioneer is that our understanding of conditioning and associative learning has gone well beyond his. It could not have done so without his paradigm. I want to talk of two rather different ways in which modern conditioning theory has progressed since Pavlov's time.

Variable Associability

Probably the most significant development in the theory of conditioning in the past 30-40 years has been the recognition of the importance of cue competition—the observation that the strength of conditioning to a CS is a function not only of its own relation to the US, but also of the relation between the US and other CSs present at the same time. The standard example of this is the blocking effect reported by Kamin (1969): if animals receive conditioning trials to a compound CS, AB, prior conditioning to A alone reduces or blocks conditioning to B. But overshadowing provides an equally important example: if A is more salient than B, relatively little conditioning will accrue to B on AB trials even if no prior conditioning was given to A alone. As I have noted, of course, overshadowing was first observed by Pavlov.

The most widely accepted account of blocking and overshadowing is that provided by the error correction model of Rescorla and Wagner (1972): the change in the associative value of a CS on any trial is a function of the discrepancy between the outcome of that trial and the associative value of all stimuli present on that trial. But other theories have suggested that cue competition is a consequence of variations in the associability of, or attention paid to, different components of a compound CS (Mackintosh, 1975; Pearce & Hall, 1980). Pavlov's own explanation of overshadowing, as his choice of terminology implies, seems to have been an attentional one—although rather different from that suggested by Mackintosh or Pearce and Hall. He reported (Pavlov, 1927, p. 141) that when a compound of a thermal and a tactile stimulus served as the CS for an acid US, the tactile stimulus alone elicited the same strength CR as the compound, but the thermal stimulus alone elicited no CR. "It is obvious, however," he continues, "that the ineffective component … could easily be made to acquire powerful conditioned properties by independent reinforcement outside the compound." (Purists may question whether this is an adequate experimental design).

Pavlov states (1927) that the mechanism underlying this overshadowing effect is "most probably a form of inhibition" (p.144), but although he continues: "This matter will be examined in detail in a subsequent lecture," I cannot find such an examination. So at this point I speculate. I *assume* that what Pavlov meant here was a form of external inhibition—where a strong stimulus presented alongside an established CS elicits an orienting reflex or attentional response, which thus distracts from the CS and prevents the appearance of a normal CR. To explain overshadowing, Pavlov must assume that the orienting reflex elicited by the stronger component of a compound CS prevents the weaker component from eliciting such a response, and that no learning of the association between a CS and US will occur in the absence of an orienting reflex to that CS. If that is correct, Pavlov's explanation of overshadowing appeals to the idea that stimuli compete for attention and that attention is necessary for learning. But this is quite different from the attentional explanations proposed by Mackintosh (1975) or Pearce and Hall (1980), where the emphasis is on *changes* in attention as a consequence of learning. In different ways, both of these last two accounts assume that the presence of the more salient component of the compound CS results in a more rapid decline in the attention paid to the less salient component (and thus what is learned about it) than would have occurred if it had been conditioned on its own.

There is no mention of changes in attention, of learning to attend to, or ignore, stimuli in the account I have attributed to Pavlov: the orienting or investigatory reflex is an inborn reflex elicited by strong, novel stimuli. However, if it is only novel stimuli that elicit the orienting reflex, then it must be affected by experience, since repeated presentation of the same stimulus will result in its habituation. We might, then, be able to credit Pavlov with an explanation for a finding of which he was wholly ignorant of, but which has seemed to many to be the prime example of changes in attention or associability—the phenomenon of latent inhibition (Lubow, 1973). Novel stimuli condition rapidly, but repeated unreinforced presentation of a stimulus before the start of conditioning trials results in relatively slow conditioning to that stimulus. Whether or not latent inhibition can be attributed to the habituation of an orienting reflex, what is lacking in the account I am attributing to Pavlov is any detailed account of the mechanism underlying such habituation. There is still no consensus, but no shortage today of theories designed to account for the decline in associability of a repeatedly presented stimulus that is observed in experiments on latent inhibition (e.g., McLaren & Mackintosh, 2000; Pearce & Hall, 1980; Wagner, 1981). What I want to go on to argue, however, is that the phenomenon of latent inhibition does not exhaust the reasons for needing to talk about changes in attention or variable associability. A second of Pavlov's discoveries, referred to above, provides evidence of a second, quite different mechanism—changes in attention to relevant and irrelevant stimuli driven by differential reinforcement.

Pavlov (1927, p.122) reported that, even after prolonged training, a dog failed to learn a discrimination between a circle and a nearly circular ellipse (with an 8:9 ratio of short to long axes). When the dog was then trained with the circle and more elliptical ellipses, however, before being tested again on the previously insoluble discrimination, on this final test he now performed perfectly accurately. Purists will again have no difficulty in questioning this experimental design: as in the case of the overshadowing experiment, the problem is that the proper control requires a between-subject design, and Pavlov's experiments were always on single animals. Nevertheless, here as elsewhere, Pavlov's findings are readily reproducible. Beginning with Lawrence (1952), numerous experiments have established that animals trained on an easy discrimination (e.g., black vs. white, green vs. yellow) and then shifted to a hard version of this discrimination (e.g., dark gray vs. light gray, green-yellow vs. yellow-green) will perform substantially more accurately on the hard discrimination than animals trained on the hard discrimination from the outset (e.g., Haberlandt, 1971; Logan, 1966; Mackintosh & Little, 1970). Like most interesting psychological phenomena, transfer along a continuum admits of several different explanations, although I can find none in Pavlov's own writings. Lawrence (1952) himself argued that training on an easy black-white discrimination made it easier for animals to learn to attend to the relevant stimuli or stimulus dimension than did training on a hard discrimination, where many more errors will be made. Any theory that says that changes in attention are driven by differential reinforcement will have no difficulty in predicting this outcome. The best evidence for this account is provided by Mackintosh and Little (1970), who found that pigeons trained on an easy wavelength discrimination performed more accurately on a hard wavelength discrimination than those trained on the hard discrimination from the outset—even if the easy discrimination was initially trained with *reversed* reinforcement assignments. That is to say: birds initially trained with yellow reinforced and green unreinforced rapidly came to perform more accurately when shifted to the hard discrimination with green-yellow reinforced and yellow-green unreinforced than birds trained on this hard discrimination from the outset. It is difficult to see how this can be explained except by saying that such easy training strengthened attention to the relevant wavelength stimuli.

My argument, therefore, is that latent inhibition and transfer along a continuum are evidence for quite different mechanisms of changes in attention, one driven by simple exposure to a stimulus, the other by differential reinforcement. An experiment on flavor aversion learning by Scahill and Mackintosh (2003) supports this argument. The design is shown in Table 1. In stage 2, all animals are trained on a difficult discrimination between two compound flavors, labeled AX and BX, where the size of the letters indicates the salience of the component flavors. The discrimination is hard because the flavor common to the two compounds, X, is salient, while the unique flavors, A and B, which differentiate the two compounds, are less salient. It follows that there are two ways of making the discrimination easier—either by increasing the salience of the two unique flavors, so the discrimination is between AX and BX, or by decreasing the salience of the common flavor, here by eliminating it altogether, so the discrimination is between A and B alone. Rats did indeed find these two discriminations substantially easier than the hard discrimination, and when shifted to the hard discrimination in stage 2, performed significantly more accurately than rats trained on the hard discrimination from the outset. Transfer along a continuum was observed.

The reason for using a flavor aversion paradigm was that it is well established that simple unreinforced exposure to two compound flavors, AX and BX, significantly facilitates the subsequent learning of a discrimination between them, and that one reason for this so-called perceptual learning effect is that such exposure results in more latent inhibition of the common X element, which is present on every exposure trial, than of the unique A and B elements, each of which is present on only half the exposure trials (see McLaren & Mackintosh, 2000, for a review of this evidence). We might expect, therefore, that stage 2 learning of the AX – BX discrimination in this experiment would also be facilitated if, in stage 1, animals simply received unreinforced exposure to the various flavors. The critical question was whether we would see an "exposure along a continuum" effect analogous to transfer along a continuum, that is to say, whether unreinforced exposure to the stimuli of the easy discriminations, AX and BX, or A and B alone, would facilitate subsequent learning of the hard AX - BX discrimination more than exposure to AX and BX themselves. It did not. Unreinforced exposure to AX and BX, or to AX and BX, did indeed facilitate stage 2 discrimination learning—but it did

Table 1

Design of an Experiment on Transfer along a Continuum in Flavor Aversion Conditioning

Groups	Stage 1	Stage 2
Easy (1)	AX vs. BX	
Easy (2)	A vs. B	AX vs. BX
Hard	AX vs. BX	

Note. A, B and X are flavors, saline, sucrose and lemon, respectively, whose salience or concentration is represented by the size of the letter.

so equally in both groups, while unreinforced exposure to A and B alone was significantly less effective—indeed did not have any effect on subsequent discrimination learning. What we found, therefore, was a clear dissociation between the effects of differentially reinforced discrimination training in stage 1, and mere unreinforced exposure. The former effect is readily explained by saying that easy discrimination training made it easier for animals to learn to attend to the relevant features of the compound flavors; the latter by saying that unreinforced exposure facilitated subsequent discrimination learning to the extent that it resulted in latent inhibition of the common X flavor. Exposure to AX and BX, or to AX and BX would result in the same level of latent inhibition to X; exposure to A and B alone failed to facilitate subsequent discrimination learning because it involved no exposure to X at all. The dissociation therefore implies that we need to distinguish two quite different processes, neither of which was anticipated by Pavlov, underlying changes in attention or variable associability. The study of simple Pavlovian conditioning has not stood still.

Associative Learning in Humans

Pavlov undoubtedly believed that the study of conditioned reflexes would advance our understanding of human behavior. But when he applied his ideas to the human condition, his application was often frankly crude: witness his lectures on "The reflex of purpose" and "The reflex of freedom" (Pavlov, 1928). The reflex of purpose is characterized as "the fundamental form of the life energy of us all," while its absence, or more properly its inhibition, is characteristic of the suicide (p. 279). Contemporary associative learning theory is rather more sophisticated in the ways in which it seeks to apply associative principles, as revealed in experiments on Pavlovian conditioning in animals, to human behavior. Note that I did not say the study of Pavlovian conditioning in humans, for it is, I think, an open question how far GSR or eyelid conditioning in people is mediated by simple associative processes, and how far more complex cognitive operations take over control of people's behavior in such experiments. The same could be said about experiments on causal or contingency judgments. Shanks and Dickinson (1988) pioneered studies that show important parallels between the behavior of people in such experiments and of animals in Pavlovian conditioning experiments, but cognitive theories of contingency judgments still flourish (e.g., Cheng, 1997).

It is not an easy matter to decide between associative and cognitive accounts of human behavior. One strategy that we have pursued is as follows. We start with an experimental paradigm where the behavior of people normally differs in an important respect from that of, say, pigeons, and where it is plausible to attribute that difference to the operation, in people, of some cognitive process absent in pigeons. We then change the experimental procedure in

a way calculated to make it harder for people to use that cognitive operation. If the effect of this manipulation is to eliminate the difference in the behavior of people and pigeons, we conclude that their now similar behavior can be explained by application of the same simple associative theory. I illustrate this strategy with some experiments on relational learning and the peak shift.

Pigeons are adept at learning, and remembering, large numbers of visual discrimination problems. But although there have been heroic efforts to prove that they are capable of learning generalized matching or same-different discriminations (e.g., Cook, Cavoto, & Cavoto, 1995; Wright, Cook, Rivera, Sands, & Delius, 1988), the most striking impression one receives from these studies is how astonishingly difficult pigeons find it to learn about the *relationship* between two or more stimuli (Mackintosh, 2000). They do not even show convincing evidence of learning about such physical relationships as "taller than" or "brighter than," when trained on a simple height or brightness discrimination (Pearce, 1991; Wills & Mackintosh, 1999). What they will show, following such training, is a peak shift —a phenomenon first discovered by Hanson (1959), but which is illustrated here in Figure 1. Hanson trained pigeons on a wavelength discrimination, with light of 550 nm as S+ and light of 560 nm as S−. When tested in extinction with wavelengths ranging from 480 to 620 nm, their peak rate of responding was not to S+, but to stimuli close to S+, but away from S−, here 540 and 530 nm. Although such behavior may look like relational transposition (they were trained to respond to the shorter of two wavelengths, and on
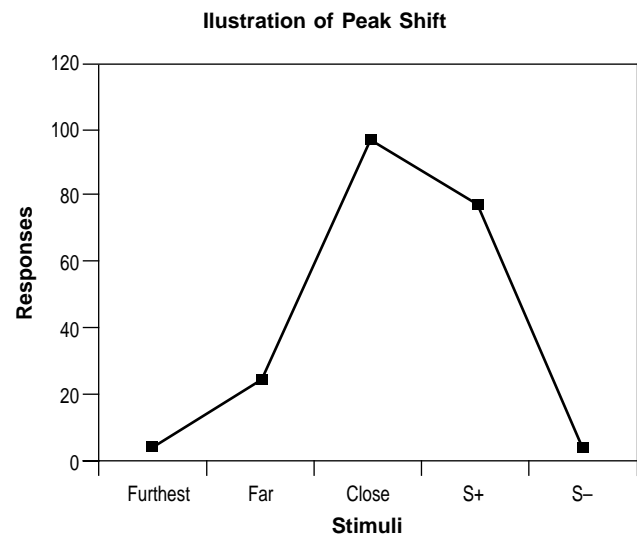

Figure 1. A schematic illustration of the peak shift—not real data, but entirely typical of the results of numerous experiments following Hanson's (1959) original demonstration. Subjects respond more to S+ than to S−, but even more to a novel test stimulus close to S+ on the opposite side from S−. When they are tested with stimuli far from S+, and with the furthest away stimulus, however, responding declines sharply.

tests responded even more to even shorter wavelengths than to their S+), we know that it is more economically and accurately explained by appeal to interacting gradients of excitation and inhibition conditioned to S+ and S−. This was first shown by Spence (1937), but Table 2 provides an elemental illustration of the argument. Stimuli are conceptualized as sets of elements, with each element of S+ and S− entering into association with the outcome of each training trial. The net associative value of S+, S−, and three test stimuli (close to, far from, and furthest from, S+) is also shown, and as can be seen, although S+ has greater value than S−, so the discrimination is learned, it has *less* than the close stimulus, so a peak shift is shown. The best evidence for this analysis is the observation that the pigeons' responding declines sharply when they are tested with even shorter wavelengths, i.e., stimuli far from S+ (in Hanson's experiment, 520 and 510 nm). As can be seen in Table 2, this is exactly the result predicted by Spence's analysis, or its elemental version.

It is not, however, how people behave. Wills and Mackintosh (1998) trained both people and pigeons on comparable brightness discriminations, and tested them with stimuli even brighter, or darker than, the stimuli on which they were trained. Pigeons showed a standard peak shift: they responded more to a test stimulus close to S+ than to S+ itself, but responding fell off when they were tested with a stimulus even further removed from S+. People, on the other hand, showed no such decline in performance to "far" stimuli: the brighter (or darker) the test stimulus, the more accurately they responded. The implication, confirmed by interviewing participants after the experiment, is that they had solved their initial discrimination relationally—that is to say, they had learned to make one response to bright, or brighter, stimuli, and another to dark, or darker, ones. Aitken (1966; see Mackintosh, 1997) has shown a similar pattern of results when pigeons and people are trained on a categorization problem—to make one response to short and fat wedge-shaped stimuli, and another to long and thin wedges. Here too, pigeons showed a peak shift: they responded more accurately to stimuli somewhat shorter and

fatter, or longer and thinner, than the stimuli on which they had been trained; but as the stimuli moved even further away from the training set, so the accuracy of their performance declined. People, however, responded more accurately and more rapidly, the longer and thinner, or shorter and fatter, the test stimulus was.

A plausible assumption is that such relational responding depends on people's ability to label the stimuli on which they are trained, as bright or dark, or short and fat or long and thin. One way to prevent such relational learning, then, would be to use stimuli, and a stimulus dimension, to which no labels could be attached. In an attempt to achieve this, Wills and Mackintosh (1998) constructed stimuli from a set of wholly arbitrary elements or "icons," in a manner similar to the stimuli illustrated in Table 2. The stimulus dimension was constructed, as in Table 2, by removing some elements or icons from each stimulus, and adding new ones, as one moved along the dimension. That this did generate a dimension is shown by the fact that, having learned to discriminate between two neighboring stimuli, pigeons showed a standard peak shift when tested with stimuli close to, and further away from, their initial S+ and S−. That people were unable to label these stimuli is shown by the fact that they too now showed a peak shift, responding more accurately to stimuli close to their S+ and S−, but less accurately to "far" stimuli.

In a similar manner, Aitken (1996) was also able to show a true peak shift in people trained on the short-fat, long-thin categorization task, by turning it into something like an implicit learning problem. Participants were required to respond as rapidly and accurately as possible to a target stimulus that appeared on the left or right of a computer screen. On some trials, the location of the target was in fact predicted by the shape of the wedge that appeared on the screen, along with a lot of other irrelevant stimuli, at the beginning of the trial (on most trials, no wedge appeared, and there was no way of predicting on which side the target would appear). Although very few participants explicitly realized this, they must all have learned it, since they all had faster reaction times on trials when the predictive wedge

Table 2

Schematic Illustration of an Elemental Analysis of the Peak Shift

	Stimuli				
	1 (furthest)	2 (far)	3 (close)	4 (S+)	5 (S−)
Elements	0112	1223	2334	3445	4556
+		+	+++	++++	+++
−			−	− − −	− − − −
Net +/−	0	+ 1	+ 2	+ 1	− 1

Note. Stimuli are conceptualized as sets of elements, here labeled 1 to 6, with each stimulus having a "central" element plus two flanking elements. Each element of S+ and S− enters into association with the outcome of each training trial, so elements common to S+ and S− are associated with both + and − outcomes. The net associative value of each stimulus on the dimension is given by the difference between + and − for each element.

appeared than on trials when it did not. When now tested with even shorter and fatter or longer and thinner wedges, they showed a standard peak shift, responding faster and more accurately to "close" test stimuli, but slower and less accurately to "far" stimuli. Interestingly enough, the two or three participants who *did* realize that there was a relationship between the shape of the wedge and the side on which the target stimulus appeared did not show a peak shift: they responded even more accurately to "far" stimuli. Our interpretation of these results is that, by embedding the predictive relationship between wedge and target in a lot of irrelevant information, and requiring participants to respond as rapidly as possible, we made it very difficult for them to articulate this relationship. Nevertheless, their associative system detected, and learned about, the contingency between wedge and target. Associative learning is all about detecting contingencies.

Conclusion

These last two experiments show that there are circumstances where people will solve discrimination or categorization problems in much the same way as pigeons, and in a manner entirely consistent with an associative analysis. That there are circumstances where they do not – because, for example, they detect relationships between stimuli—is not very surprising. The importance of these results, I venture to suggest, is that they reveal the operation of a simple associative learning system in people, when steps are taken to make it harder for them to deploy the cognitive resources normally available to them.

This associative learning system, as Pavlov certainly believed, is indeed of widespread importance—common to most if not all vertebrates, and also to some invertebrates. It will solve a wide variety of problems, many superficially quite complex, and it is what enables other animals, and ourselves, to gain an understanding of the causal structure of the world. Pavlov's great insight was to see that it could be successfully studied in a very simple preparation, and that the development and exploitation of the paradigm he devised was the critical step needed to generate a coherent body of knowledge about associative learning, which in turn was the critical step needed to develop an adequate theory of associative learning. If today we have made some progress towards this latter goal, that is thanks to Pavlov's pioneering work.

References

Aitken, M. R. F. (1996). *Peak shift in pigeons and human categorisation.* Unpublished doctoral dissertation, University of Cambridge.

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review, 104*, 367-405.

Cook, R. G., Cavoto, K. K., & Cavoto, B. R. (1995). Same-different texture discrimination and concept learning by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 21*, 253-260.

Guthrie, E. R., & Horton, G. P. (1946). *Cats in a puzzle box.* New York: Rinehart Press.

Haberlandt, K. (1971). Transfer along a continuum in classical conditioning. *Learning and Motivation, 2,* 164-172.

Hanson, H. M. (1959). Effect of discrimination training on stimulus generalization. *Journal of Experimental Psychology, 58,* 321-334.

Kamin, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42-64). Halifax: Dalhousie University Press.

Konorski, J. (1948). *Conditioned reflexes and neuron organisation.* Cambridge, UK: Cambridge University Press.

Lawrence, D. H. (1952). The transfer of a discrimination along a continuum. *Journal of Comparative and Physiological Psychology, 45*, 511-516.

Logan, F. A. (1966). Transfer of discrimination. *Journal of Experimental Psychology, 71,* 616-618.

Lubow, R. E. (1973). Latent inhibition. *Psychological Bulletin, 79,* 398-407.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82*, 276-298.

Mackintosh, N. J. (1995). Categorisation by people and pigeons. *Quarterly Journal of Experimental Psychology, 48B*, 193-214.

Mackintosh, N. J. (1997). Has the wheel turned full circle? Fifty years of learning theory, 1946-1996. *Quarterly Journal of Experimental Psychology, 50A*, 879-898.

Mackintosh, N.J. (2000). Abstraction and discrimination. In C.Heyes & L.Huber (Eds.), *The evolution of cognition* (pp. 123-141). Cambridge, MA: MIT Press.

Mackintosh, N. J., & Little, L. (1970). An analysis of transfer along a continuum. *Canadian Journal of Psychology, 24,* 362-369.

McLaren, I. P. L., & Mackintosh, N. J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning and Behavior, 26,* 211-243.

Moore, B. R., & Stuttard, S. (1979). Dr. Guthrie and Felis domesticus, or, tripping over the cat. *Science, 205,* 1031-1033.

Pavlov, I. P. (1927*). Conditioned reflexes.* London: Clarendon Press.

Pavlov, I. P. (1928). *Lectures on conditioned reflexes.* New York: International.

Pavlov, I. P. (1932). The reply of a physiologist to psychologists. *Psychological Review, 39,* 91-127.

Pearce, J. M. (1991). The acquisition of abstract and concrete categories by pigeons. In L. Dachowski & C. Flaherty (Eds.), *Current topics in animal learning: Brain, emotion and cognition* (pp. 141-164). Hillsdale, NJ: Erlbaum.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87,* 532-552.

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

Rosenzweig, M. R. (1962). The mechanisms of hunger and thirst. In L. Postman (Ed.), *Psychology in the making* (pp. 73-143). New York: Alfred A. Knopf.

Scahill, V., & Mackintosh, N. J. (2003). Transfer along a continuum and perceptual learning in flavour aversion conditioning (submitted).

Shanks, D. R., & Dickinson, A. (1988). Associative accounts of causality judgment. In G. H. Bower (Ed.), *The psychology of learning and motivation, 21* (pp. 229-261). New York: Academic Press.

Shaw, G. B. (1944). *Everybody's political what's what?* London: Constable.

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis.* Englewood Cliffs, NJ: Prentice Hall.

Spence, K. W. (1937). The differential response to stimuli varying within a single dimension. *Psychological Review, 44*, 430-444.

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5-47). Hillsdale, NJ: Erlbaum.

Wills, S. J., & Mackintosh, N. J. (1998). Peak shift on an artificial dimension. *Quarterly Journal of Experimental Psychology, 51B,* 1-31.

Wills, S. J., & Mackintosh, N. J. (1999). Relational learning in pigeons? *Quarterly Journal of Experimental Psychology, 52B,* 31-52.

Wright, A. A., Cook, R. G., Rivera, J. J., Sands, S. F., & Delius, J. D. (1988). Concept learning by pigeons: Matching-to-sample with trial-unique video picture stimuli. *Animal Learning and Behavior, 16,* 436-444.