eScholarship
International Journal of Comparative Psychology

Title
At the Interface of Learning and Cognition:An Associative Learning Perspective

Peer reviewed

At the Interface of Learning and Cognition:
An Associative Learning Perspective

Andrew R. Delamater
Brooklyn College of the City University of New York, U. S. A.

This paper reviews some of the literature on Pavlovian and instrumental conditioning as they relate to "cognitive" factors in behavior. Studies of Pavlovian learning have centered around the notion that a representation of the unconditioned stimulus plays a critical role in performance. However, much work will need to go into characterizing the nature of the representations that mediate learning. In particular, current research illustrates that "images" and "expectancies" of reward may differ in fundamental ways, and also that learning about temporal, motivational, and sensory properties of reward might involve different systems. The study of instrumental learning also poses challenges for addressing the question of what representations, i.e., associative structures, underlie such learning. Current work reveals a host of associative structures that may participate in learning and performance though how these different structures participate in a unified approach is currently unknown. The associative approach can be contrasted with inferential reasoning approaches to instrumental action, and there are two key findings that seem outside the scope of a reasoning approach. Nevertheless, future work will be required to determine just how far purely associative models will be able to go in order to account for complex behavior.

It is quite remarkable how psychology has evolved to regard basic "learning" and "cognitive" processes in such distinct ways. Part of the blame, perhaps, can be traced to the radical behaviorism of the middle part of the 20[th] century that so strongly dismissed "mental" or "internal" constructs as having anything to do with the study of behavior. Although there were successful attempts at bringing the internal into the study of behavior (e.g., Tolman, 1932), indeed, to the forefront of the study of behavior within the field of learning theory, the momentum of "The Cognitive Revolution" by the late 20[th] century firmly took hold, more generally, on psychology. Unfortunately, this meant that any remnant of a "behavioral" approach to psychology became equated with its most radical forms. Important distinctions between behavior theories emphasizing control by internal constructs (like "purpose" and "expectation") and those that were restricted to observables were lost by most who were enthusiastic with the new information processing framework (Atkinson & Shiffrin, 1971) that has so dominated the study of cognition ever since. The consequence of this is that in most psychology departments what has become the study of animal learning, to the extent that it can be found at all, has been regarded by most as peripheral, antiquated, or too low-level to have anything to do with the study of human cognition. That stuff is merely "associative" seems to be the prevailing view.

I believe that the reality of the situation, however, is that the associative model provides us with the most biologically plausible way of understanding how the brain works to provide us with thoughts, feelings, motivations, decisions, actions, and the like. In my view, a major development in associative theorizing

occurred starting in the late 1980s when "parallel distributed processing" and "connectionist" modeling efforts began to show how complex human cognitive functions can be understood to emerge from very simple associative systems whose basic operating principles included constructs no more complex than "excitation," "inhibition," and "associative strength," the very essence of associative theory ever since its inception (see O'Reilly & Munakata, 2000; Rumelhart, McClelland, & the PDP research group, 1986).

Nevertheless, although the success of connectionist models to explain aspects of more complex cognition is undeniable, I believe that a prevailing view is that such explanations are at a different level of analysis that need not really concern us when trying to explain psychological processes at a higher level of analysis. In other words, while the connectionist framework may be good for implementing explanatory constructs at the psychological level to a lower neural level of analysis, they need not replace those explanations at the higher psychological level. This prevailing view is unfortunate, however, because it means that when perfectly adequate explanations exist at an associative level, theorists might be disinclined to apply those explanations in favor of ones that may invoke constructs considerably more complex than those assumed within the associative framework. The issue here is one of parsimony.

But really what makes it so startling to realize that psychologists have evolved to think so differently about "learning" and "cognitive" processes is that there seems to be an assumption that what is cognitive could not possibly be associative in nature. Although there may be historical reasons, as noted above, for understanding this, and there may be theoretical levels of analysis issues to contemplate as well, it is, perhaps, wiser for us to attempt to tackle head-on the basic issue of what the difference might be if there is one. On the one hand, most learning theorists, I believe, would accept the view that behavior is complexly governed by the interaction of motivation, representation, memory, and decision processes. On the other hand, most cognitive theorists would also probably agree that motivation, representation, memory, and decision processes govern behavior in some complexly determined way. The difference is that the cognitive theorist will add to the mix "inferential reasoning" and/or "rational inference" processes, while most animal learning theorists will not. Learning theories will usually attempt to explain the complex interactions assumed here within some sort of, essentially, reflexive system, i.e., an associative system. Some have referred to this as the "excitatory links" model (e.g., Dickinson, 1989). Within such a framework it is hard to see how concepts like inferential reasoning or rational inference can be imagined. It is important to realize, however, that a system that is essentially associative in nature, ultimately, may perform in such a way to produce behavior that appears "rational" or that obeys rules of "inferential reasoning".

This approach would regard inferential reasoning, though, as an emergent property of the operation of an underlying connectionist system, and not as some attribute of the system that falls outside of the scope of an associative process. It may, of course, be quite fruitful for the researcher to work at this more emergent level of analysis in order to learn more about the empirical properties of the system. However, a complete understanding will not be reached until it can be understood how these emergent properties arise out of an essentially

associationistic system to begin with, since the underlying neurobiology of the system is, in fact, connectionistic. In this sense, the "cognitive" level analysis can help inform the "associative" level of analysis about the global properties of the system and this can be helpful in the construction of more complete models. At the same time, when the "associative" level of analysis can provide adequate explanations of the "cognitive" level phenomena then the law of parsimony would dictate that the particular mechanisms involved need be no more complex that are assumed in the associative-level account. In this scenario, the two approaches are entirely complementary, not antagonistic. It would surely be interesting if one could derive predictions about behavior based on a rational, inferential processing model that could not also be anticipated from the perspective of an associative theory. Here is where working at the "cognitive" level might prove to be rather fruitful. If this could occur, however, it would not require that we abandon the idea that behavior is controlled through associative mechanisms. It could merely point to some way in which our current associative models would require modification. Exactly this sort of controversy has been stimulated most recently by Blaisdell and his colleagues in their work on "causal reasoning" in rats (Blaisdell, Sawa, Leising, & Waldmann, 2006).

        The purpose of the present article is not to reach any resolution to these important considerations of "learning" or "cognitive" approaches. Rather, I hope to show how successful the associative framework has been at providing us with information about some key issues that are rather "cognitive" in their feel. The study of Pavlovian conditioning has been strongly guided by the assumption that learning consists of a connection between mental representations of events (e.g., Pearce & Hall, 1980). Indeed, some have argued that learning of an associative relationship between two events, conditioned and unconditioned stimuli (CS, US respectively), can be thought of as establishing a memory structure that is essentially "declarative" in its content (e.g., Dickinson, 1980). In this way, one can think of the study of associative content as the study of the structure of knowledge. This approach is really no different from that advocated by the early empiricist philosophers who took associative learning as a means of understanding mind (e.g., Boakes, 1984). In more recent years, a number of interesting discoveries have been made concerning the study of associative content in Pavlovian conditioning, and the purpose here will be to review some of the key findings. This will hopefully illustrate how an associative approach can lead to a quite rich set of knowledge structures in the brains of organisms. That these knowledge structures play a key role in determining action and action selection will then be briefly reviewed in another section that attempts to identify some of the key associative structures that are thought to govern the control of instrumental behavior. While illustrating this I will also point out some places where the control of instrumental action appears to be better described from the perspective of an associative framework rather than an inferential reasoning perspective. Finally, I will return to the issue of causal reasoning in rats and briefly discuss one way in which an associative approach might address some of the key findings. My overall aim is to show the richness of the associative approach in grappling with behavior that seems, on the face of it, to be essentially "cognitive" in its essence.

On the Nature of Reward Representations and the Structure of Knowledge

While most theorists will agree that Pavlovian conditioning consists of the establishment of a connection between mental representations of the CS and US, the precise nature of those representations is less clear. I have elsewhere discussed the possibility that CS representations are dynamic structures that might change across a conditioning phase (e.g., Delamater, 1998; 2011; Delamater, Sosa, & Katz, 1999). Another important consideration for which there is a large database is the question regarding the nature of the US representation. At least since Konorski (1967) there has been widespread agreement that the brain represents the US in terms of its motivational/emotional impact on the organism as well as its more sensory-specific qualities (e.g., Mackintosh, 1983; Wagner & Brandon, 1989). When a CS is paired with a US it may, in turn, come to form separate associations with both of these qualities of the US.

However, the US is surely a more complex event than is captured by these ideas. The US is not only a biologically important event with emotional/motivational and sensory qualities. It is an event that occurs at a specific point in time, with reference to the onset of some predictive cue. It is also an event containing short-lasting hedonic qualities. And further, it is an event that is itself capable of evoking specific responses unconditionally. We have earlier asked to what extent the CS might enter into associations with each and every one of these qualities of the US, and to what extent the learning of these multiple associations might depend upon truly independent systems (see Delamater & Oakeshott, 2007). Although we do not yet have answers to all of these important questions, the study of learning from the point of view of identifying the nature of the US representations that enter into learning has provided us with some interesting discoveries.

Images and Expectancies

There is an interesting distinction that is now emerging in the study of US representations in Pavlovian learning between images and expectancies. While a large amount of data exists to show that when CS and US are paired in such a way so as to support learning, the CS enters into an association with the sensory-specific qualities of the US. The most simple way of demonstrating this is to show that conditioned responding to the CS is reduced when it is tested after the US has been devalued through some treatment following the conditioning phase and independently of the CS. This "US devaluation effect" has been well documented and points to the conclusion that the CS must have associated with some specific sensory property of the US that itself was diminished in value (for an early review see Delamater & LoLordo, 1991). The CS is said to evoke a specific representation of the US, or encodes some highly specific sensory quality of the US (e.g., Colwill & Motzkin, 1994; Holland, 1990).

However the exact nature of this specific US representation is less clear from these studies alone. Hall (1996) earlier speculated that the CS might come to associatively activate a representation of the US that is the very same

representation as that actually activated by the US (or at least significantly overlaps with it). We now have some evidence to suggest that this is partly true.

Delamater, LoLordo, and Berridge (1986) some time ago demonstrated that rats will display the same set of orofacial taste reactivity responses to water in the presence of distinct auditory cues for either a highly palatable sucrose solution or a highly unpalatable quinine solution as they display to sucrose and quinine themselves. This result is suggestive of the view that the rats are, in effect, experiencing the water as though it was sucrose in the presence of the sucrose cue and quinine in the presence of the quinine cue. However, although this finding is suggestive, it is not definitive.

More recently, in a series of studies Peter Holland and his colleagues have performed a number of studies that more clearly make the case. Kerfoot, Agarwall, Lee, and Holland (2007) used the same sort of procedure as Delamater et al. (1986), but they also examined the effects of US devaluation. Initially rats were given pairings of a tone with liquid sucrose. Control rats were given these events but unpaired. Then, separate groups of rats received sucrose and LiCl either paired or unpaired to form four groups in total. This was designed to reduce (or maintain) the value of the taste of sucrose. Finally, all rats were tested with presentations of tone followed by plain water, and their licking behavior was closely videotaped and monitored. Just as Delamater et al. (1986) had observed, rats that received tone-sucrose pairings and did not receive sucrose devaluation displayed more appetitive licking behaviors to water than control animals. However, rats that received tone-sucrose pairing followed by sucrose devaluation displayed low levels of appetitive licking behaviors and they further displayed increased levels of aversive licking behaviors than any of the other groups. The results of this study show more clearly that the CS evokes a representation of sucrose that is quite similar to that actually evoked by sucrose itself, as though the rats actually perceive sucrose when water is presented in the presence of the cue for sucrose.

These authors went on to demonstrate that individual cells within the nucleus accumbens shell were differentially activated by the tone + water combination depending upon whether sucrose had been devalued or not. Earlier, Berridge and his colleagues provided evidence for different "hedonic hot spots" within the anterior and posterior regions of the accumbens shell (e.g., Reynolds & Berridge, 2002). That work demonstrated that palatable solutions more strongly activated cells within the anterior region of the accumbens shell, whereas unpalatable solutions more strongly activated cells within the posterior region. Kerfoot et al. (2007) demonstrated with cFos staining techniques that the sucrose-paired tone together with water evoked more cells in the anterior region of the shell if sucrose had not been devalued, but more cells in the posterior region of the shell if sucrose had been devalued. These results are highly supportive of the view that the tone causes the rats to treat water as though it actually was sucrose.

Further work by Saddoris, Holland, and Gallagher (2009) suggest that the gustatory cortex may also participate in this learning. Using additional immunohistochemistry techniques these authors demonstrated that a cue of sucrose and sucrose itself activated the very same cells within the gustatory cortex at levels greater than that seen in unpaired controls. Putting these facts together, it appears that the CS might activate a specific image of sucrose that alters the manner in

which the rat perceives water. Similar findings were reported recently by Desgranges, Ramirez-Amaya, Ricaño-Cornejo, Levy, and Ferreira (2010) who found larger numbers of cells in the basolateral amygdala that were responsive to both a cue for sucrose and sucrose itself than in unpaired controls.

Other work, however, suggests that this conclusion is an oversimplification. Holland, Lasseter, and Agarwal (2008) demonstrated that while a sucrose-paired tone had the effect of changing the way the rat responds to plain water, the tone loses this ability after extensive training (see also Holland, 1998). In this study, tone was paired with sucrose for either 16 or 112 pairings in different groups. Different subgroups then either received sucrose-LiCl devaluation training or unpairings of these events. All groups then were given tone + Water test trials and licking behavior was observed. It was found that after minimal tone-sucrose training the tone affected water licking in a manner that was dependent upon whether sucrose was devalued or not – more ingestive licks if sucrose was valued than if it was devalued. However, this difference in licking was strongly attenuated in rats given extensive tone-sucrose training to begin with. It is as if the ability of the tone to activate a specific image of sucrose is diminished with overtraining.

However, these authors also provided another measure of learning – conditioned magazine approach responding. With this measure of learning the rats given 16 and 112 pairings, alike, both reduced their magazine approach responses after sucrose had been devalued compared to when sucrose had not been devalued. This classic US devaluation effect shows that the tone evoked a representation of sucrose that was specific in its sensory content equally well in animals trained with 16 or 112 pairings (see also Holland, 1998). How are we to reconcile the different results with these different measures? One answer is that with minimal training the CS evokes a representation of the US that is indistinguishable from the US, but with over training the CS evokes a representation of the US that preserves at least some of its sensory qualities but which is distinguishable from the US itself. The first representation can be referred to as an "image" of the US while the second might be referred to as an "expectancy" of the US. In other words, the rat seems to perceive water as though it was sucrose in the presence of the sucrose cue following limited training, but after extensive training the rat seems to understand that the water is not sucrose but that sucrose is, nevertheless, expected to occur in the food magazine.

The results are highly provocative, but they do not address the issue of why the representation seems to undergo this transition from initial to extended training. One possibility is that with extended training subjects have a greater opportunity to contrast the image of the US with the actual US, and this results in a sort of perceptual learning effect. Alternatively, as Holland (1998) suggested there may be multiple representational systems that get recruited at different times over the course of conditioning. Although these analysis are highly speculative, overall, these studies give us great insight into the nature of the sensory-specific representations of the US that can come to guide behavior in interesting ways.

Sensory and Temporal Expectancies of the US

It has long been recognized that the time of arrival of the US plays a significant role in Pavlovian learning. Temporal contiguity effects have been observed in virtually every learning paradigm studied (see Mackintosh, 1974, 1983). Exactly how time plays a role in conditioning, however, continues to be debatable. Some authors are of the view that the appropriate framework for understanding Pavlovian learning is not the associative framework, but rather is one that emphasizes temporal learning processes (e.g., Balsam & Gallistel, 2009; Gallistel & Gibbon, 2000). Other authors suggest that timing processes can be understood in associative terms (e.g., Bouton & Garcia-Gutierrez, 2006; Bouton & Hendrix, 2011; Kehoe, Horne, Macrae, & Horne, 1993). Without attempting to solve the problem of timing, however, one can ask a somewhat simpler question, and that is whether learning to represent the time of arrival of the US depends upon a learning system that is similar to or distinct from learning about other qualities of the US (e.g., sensory).

In a series of studies Delamater and Holland (2008) examined this question using a Pavlovian to instrumental transfer test to assess learning of these different US attributes. Their results suggest that the CS does convey information both about the specific sensory and temporal properties of the US, and also that these might depend upon separable processes. In these studies rats were initially given instrumental training with two different responses each being reinforced with distinct USs (R1-US1, R2-US2). Subsequently, the rats were given Pavlovian delay conditioning with two separate CS-US pairs (CS1-US1, CS2-US2). Different groups of rats were trained with different CS-US intervals (ranging across experiments from 2 – 180 s). Finally, the rats were given an instrumental choice test under extinction conditions where each CS was presented on occasional test trials. In this test, evidence for sensory-specific learning was obtained by the fact that the rats reliably chose the instrumental response that had earlier been reinforced with the same US as that signaled by the CS (e.g., CS1-R1, CS2-R2).

Evidence for temporal learning was obtained by assessing the influence that these CSs had on instrumental choice at different moments in time within the CS. Groups trained with 30-s and 60-s CS-US intervals, for instance, displayed specific transfer that was greatest towards the end of the CS close to the time of anticipated US delivery. In other words, the specific Pavlovian-instrumental transfer effect increased over the CS-US interval and was largest around the time of anticipated US delivery. These results suggest that the CS evokes representations of the US that contain information not only about "what" the US is but also "when" it will occur.

These authors also asked if learning about these different US attributes might depend upon the same or distinct underlying systems. The results are more in support of the distinct systems view. First, they observed that subjects trained with a 180 s CS-US interval displayed reinforcer-specific Pavlovian-instrumental transfer, indicating good "what" learning. However, these subjects did not also display temporally-specific Pavlovian-instrumental transfer like subjects given training with shorter CS-US intervals had. This result suggests that expression of

"what" learning does not depend upon expression of "when" learning, and, thus, suggests a way in which the two might be independent of one another.

Second, when assessing conditioned magazine approach response timing during the training phase, it was observed that the temporal distribution of responding was sharper with shorter intervals (even after plotting the results on a relative time scale). However, it was also observed that training with the shortest CS-US interval resulted in no evidence of US-specific Pavlovian-instrumental transfer indicative of "what" learning. Moreover, other measures of sensory-specific associations (selective US devaluation and selective potentiated feeding) also failed to provide evidence that the CSs trained with a very short CS-US interval revealed evidence of "what" learning. Thus, it appears as though from these results that CS-US interval affects US timing ("when" learning) and US-specific Pavlovian-instrumental transfer ("what" learning) in different ways. This also suggests that the two forms of representations underlying Pavlovian learning may depend upon distinct systems, though more work is needed to strengthen this argument.

We have recently begun exploring this possibility further using a US preexposure protocol. Our initial study asked how first learning "when" a US would occur would subsequently affect learning "whether" the US would occur. In this study, three groups of rats were initially given magazine training with a pellet US. One group was magazine trained over two days with a fixed time 30" schedule (FT 30), a second group with a variable time 30" schedule (VT 30), and the control group was trained on one day with the FT 30 schedule and on the other day with the VT 30 schedule. For the next 19 days the FT 30 and VT 30 groups continued to receive pellets on their appropriate schedules, but the control group was placed in the experimental chambers without food pellet deliveries. The conditioning phase followed these US preexposure sessions. During this phase, all three groups of rats received light-Pellet pairings, where the CS-US interval was 30" and the intertrial interval averaged 3 min. If "when" learning plays a significant role in Pavlovian conditioning, then we might expect the FT 30 group to have learned during the preexposure phase that pellets occur after a 30-s interval, and that this learning might help them learn that the pellet will occur 30" following the onset of light. Other studies have shown that the representation of time is not specifically tied to a particular stimulus, but is rather more abstract (Meck & Church, 1982). In contrast, Group VT 30 subjects should be severely impaired compared to the minimally preexposed (and then extensively extinguished) control group because they would not have had the opportunity to learn that the pellet occurs after a specific period of time.

The results are depicted in Figure 1 and partially confirmed these predictions. We observed that the control group learned to approach the food magazine during the light CS in a temporally-specific manner more rapidly than the FT 30 group who, in turn, learned more rapidly than the VT 30 group. Figure 1A presents mean magazine approach responding in each of the groups over successive training blocks in the form of the slopes of the best fitting line (from the least squares regression equation) relating magazine approach responding across the 30-s interval between CS and US onsets. If subjects have learned to anticipate the delivery of food at the very end of the interval, then responding should be low
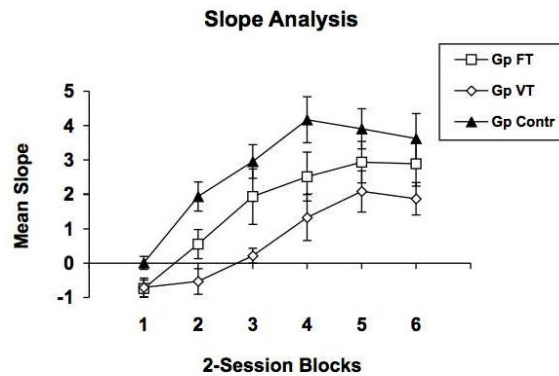
at CS onset and steadily increase around US onset. This would be reflected in a positive slope. Note how the control group rapidly acquired this pattern of responding, but the FT and, especially, the VT groups were slow. An analysis of variance applied to these data revealed a significant Group main effect, $F(2,21) = 7.17$, $p = 0.004$, as well as a significant Block main effect, $F(5,105) = 33.68$, $p = 0.0001$, but no interaction between these two factors. The graph shows clear separation among the three groups at various times throughout training.

Figure 1B shows magazine approach responding during successive 5 s bins within the CS during the 2^{nd} two-session block of training for illustrative purposes. Note that the control group developed a pattern of accelerated responding across the CS in this block indicative of good timing of the conditioned response, but that this pattern was greatly attenuated in Group VT and to a lesser extent in Group FT. An analysis of variance applied to these data revealed a significant main effect of Interval, $F(5,105) = 4.34$, $p = 0.001$, as well as a significant Group x Interval interaction, $F(10,105) = 5.65$, $p = 0.0001$, to support these impressions.
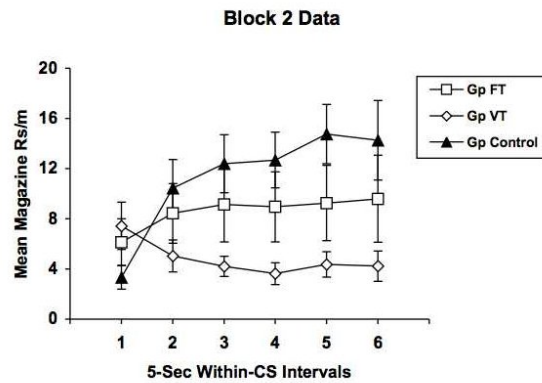
One explanation for these findings is based on the classic explanation of the US preexposure effect. First consider why group VT 30 may have shown poor learning. The classic explanation is based on the notion of context blocking (e.g., Randich & LoLordo, 1979). According to this view, during the preexposure phase VT 30 subjects learned a strong context-pellet association. This association then blocked conditioning of the light-pellet association during the training phase, according to principles of the Rescorla-Wagner model (Rescorla & Wagner, 1972). However, since the FT 30 group received the same number of pellet preexposures as the VT 30 group some explanation will have to be offered as to why the FT 30 group was more successful than the VT 30 group during light-pellet training. One possibility is that the context was less strongly conditioned during the preexposure phase in FT 30 subjects. This could have occurred if it is assumed that temporal cues came to overshadow context cues in their association with the pellet during the preexposure phase.

One problem, however, for this analysis is that when we assessed context conditioning during the training phase we saw no evidence that this was reduced in FT 30 subjects. We assessed this by looking at magazine approach responses in the pre-CS periods on the first day of training. During the first session of training the pre-CS magazine approach response rate was highest in Gp FT 30, intermediate in Gp VT 30, and lowest in the Control group (see Fig. 1C). These differences were supported by a significant main effect of group, $F(2,21) = 5.12$, $p = 0.016$. Thus, it appears from these results that context blocking is not an adequate account of our findings. Instead, it would appear as though there are two distinct forms of learning occurring here, learning "when" the US will occur and learning "whether" the US will occur. The preexposure manipulation impairs learning "whether" the US will occur, but the FT/VT difference is accountable in terms of different effects on "when" learning.
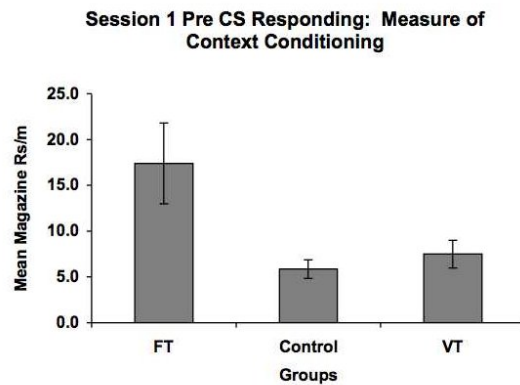
A)

Slope Analysis

B)

Block 2 Data

C)

Session 1 Pre CS Responding: Measure of Context Conditioning

Figure 1. Panel A shows the mean slope of the best fitting line for magazine approach responding across six 5-s intervals within the CS at different blocks of training for each of the three groups. Gp FT was preexposed to the pellets according to a fixed time 30 s schedule, and Gp VT was preexposed to pellets according to a variable time 30 s schedule. All three groups received Light-Pellet pairings with a fixed 30 s CS-US interval throughout training. Panel B shows magazine response rates distributed across the CS for each group on block 2 of training. Panel C displays differences in pre-CS magazine approach responding during the first day of Light-Pellet training.

- 398 -

Sensory and Motivational Representations of the US

A fair amount of work has also addressed the possibility that during Pavlovian learning the CS comes to enter into separate associations with specific sensory features of the US and also with the general motivational features of the US. The clearest demonstration of this comes from a study by Corbit and Balleine (2005). This study also used a Pavlovian-instrumental transfer task with rats but with an important difference to the procedure outlined above. Initially, rats were given instrumental training on two different response-US pairs (R1-US1, R2-US2). During the Pavlovian training phase two CSs were trained with the same USs used during instrumental training (CS1-US1, CS2-US2) and, in addition, a third CS was trained with a motivationally similar but qualitatively distinct US (CS3-US3). Finally, in the test phase each CS was presented in test sessions in which only one or the other of the responses was available. These authors observed that CS1 and CS2 selectively elevated above pre-CS baseline levels the response that had previously earned the same US as that signaled by the cue. Further, these CSs had no effect on the response that previously earned the other US. In addition, CS3 elevated both instrumental responses above baseline levels even though neither instrumental response had ever earned the US signaled by this cue. The authors interpreted these findings by assuming that CS1 and CS2 exerted selective transfer through its ability to activate specific representations of the USs, i.e., through the activation of specific US "expectancies." On the other hand, CS3 was assumed to have produced a general transfer effect through this stimulus' ability to activate a general motivational representation of the US. This general motivational effect of CS3 was reflected in a generalized increase in activity that included both responses.

Further evidence of a dissociation between these sensory and motivational effects was provided by the results of other groups in this study. One group was given pretraining lesions of the central nucleus of the amygdala (CN), and a second group was given pretraining lesions of the basolateral amygdala (BLA). The CN group displayed specific transfer in response to CS1 and CS2, but lost the general transfer effect in response to CS3. In contrast, the BLA group displayed general transfer in response to CS3 but these subjects lost their specific transfer effects in response to CS1 and CS2. These results convincingly point to a system dissociation in the Pavlovian learning processes involving these sensory-specific and general motivational features of the US.

In order to explore whether these same general and specific transfer effects can be seen to effect performance with humans in an analogue of this design, we have recently reported evidence for both specific and general forms of transfer (Nadler, Delgado, & Delamater, 2011). We are currently exploring with imaging technology the possibility that these different forms of transfer depend upon different neural substrates in the human brain as is indicated in the rodent work.

Summary

There has been much progress in understanding the nature of the US representations that mediate Pavlovian learning. All of these studies point to the importance of considering how the US might be represented by the brain and how such representations might enter into associations with predictive cues. The important question of the independence of these learning systems is only beginning to be addressed and will likely be better understood with further research directed at both the neural and behavioral/cognitive levels of analysis. For now it can be safely said, however, that any attempt at understanding the structure of "knowledge" will very likely require an understanding of the nature of the events that become associated. A CS can convey information about multiple attributes of the US. Which of these attributes become incorporated into the associative structure mediating performance will serve as the basis upon which knowledge about future events can be understood. The very idea that events can be represented in multiple ways that each provides different forms of information about a stimulus is an inherently "cognitive" way of describing the system. However, in these instances, at least, the "cognitive" descriptions can be readily reduced to a set of associative-level mechanisms.

Mechanisms of Response Control in Instrumental Learning

Reward representations have also been thought to play a critical role in instrumental learning and the control of instrumental performance. A wide variety of intricate studies have been performed that point to an extremely rich associative circuitry underlying the control of instrumental performance. At the same time theorists have also approached the problem of understanding instrumental responding from the point of view of a more common sense practical inference learning perspective. The point of the brief review in this section will be to highlight some of the important issues and to show some of the limitations of the inferential reasoning perspective as well as challenges to the associative perspective. More exhaustive summaries appear elsewhere (e.g., Balleine & Ostlund, 2007; Hall, 2002).

First let us consider the basic problem of why the animal engages in the instrumental response in the first place. This problem really only arises when one adopts a mechanistic account of instrumental performance. Other non-mechanistic accounts (e.g., Skinner, 1938) sidestep the issue by defining the problem as one of identifying functional relationships between behavior and environmental events, i.e., observables, and, thereby, avoid any consideration of the underlying mechanisms accounting for the occurrence of behavior. However, an associative account will begin by attempting to identify the associations learned in the situation and then attempt to understand why behavior results from the associative mechanisms so identified. One clear example comes from Thorndike's own analysis of instrumental learning (Thorndike, 1911). Thorndike assumed that once reinforcement occurred it strengthened an associative connection between the stimulus situation and the response that immediately preceded the reinforcement. This S-R association was assumed to account for why the animal engaged in the

response when it confronted the same stimulus situation in the future. The mechanism clearly provides an account of the origin of the response – the stimulus evokes it.

While there is reason to accept that this sort of associative mechanism is learned and does contribute to performance in various situations, there is also widespread acceptance that this analysis, by itself, is incomplete. One simple experiment that shows this was reported by Colwill and Rescorla (1985) using the same reward devaluation methods as have also been widely employed in the study of Pavlovian learning. Briefly, in this study two instrumental responses were trained in independent sessions with each response reinforced by a distinct outcome (R1-O1, R2-O2). Subsequently, one of these outcomes was presented on its own in the absence of any opportunity to engage in either instrumental response, and this was immediately followed by an injection of the emetic LiCl. On another day, the other outcome was presented without LiCl. Thus, one of the reinforcing outcomes was devalued with the value of the other maintained. Finally, the animals were given an opportunity to choose between the two instrumental responses during an extinction session without any reinforcers present. It was observed that the rats preferred the response whose associated outcome had not been devalued during this test session. The result argues against Thorndike's S-R associative account because that explanation has no way of explaining how reward devaluation could have a selective effect on responding since the reward itself is not encoded into the associative structure assumed to underlie learning.

The result is most naturally understood by assuming that the animal forms an association that reflects the actual programmed contingency, namely, an R-O association. While this sort of result is widely held to reflect such learning it raises the further question as to why the animal actually engages in the response to begin with. The problem arises when we consider that an association between two events normally only permits for activation to "flow" from a representation of the first event to the second. In this case, the rat presses the lever for food reward and then acquires an association between the lever press response and food, but why does this association result in the rat actually pressing the lever in the future? The mechanism allows for a mental representation of the response to activate a mental representation of the outcome, but what actually evokes the response? This problem is reminiscent of the famous complaint that Guthrie (1935) raised against Tolman's rats being "buried in thought," and there are two rather different types of explanation offered.

The first explanation has its origins in the writings of Pavlov (1932). In this article, Pavlov suggested that while the animal could learn an association between the response and the rewarding outcome, it could use that association in the backward direction for the purposes of controlling the response. Thus, all that needs to happen is that the outcome representation be activated, and this will, in turn, cause the response to be activated by virtue of the already established R-O association. This has come to be known as the "bidirectional" hypothesis (see also Mackintosh & Dickinson, 1979), and is one mechanism used commonly to explain selective Pavlovian-to-instrumental transfer of control. Indeed, Pavlov suggested that when one attempts to teach their dog to give its paw, the verbal command acts

as a cue to trigger the thought of the outcome, which, in turn, triggers the response through this bidirectional response-outcome association.

This approach has some merit, but another radically different approach has also been proposed. Mackintosh and Dickinson (1979) suggested that the rat represents this R-O association in propositional form, once it is established, i.e., in terms of a "belief" about the consequences of its actions. If the animal has a "belief" about how food can be obtained, and it also has a "desire" to obtain food, then it can combine these two propositions and derive the practical inference to respond. At the heart of this analysis is the view that the animal engages in a rational decision making process. Other work that I will soon review suggests that sometimes the rat responds in seemingly non-rational ways. However, the general approach has also been usefully applied to studies of "incentive learning" (Dickinson & Dawson, 1988) and "causal reasoning" (Blaisdell et al., 2006).

Incentive learning refers to the finding that before an animal can respond appropriately to a shift in motivational state, it must first have some understanding of the value of the goal object in that shifted motivational state. A clear example comes from an early study by Balleine (1992). In this study, rats were trained to press a lever for food pellets while in a hungry state. Subsequently, some of the rats were given an extinction test while hungry and others after they had been sated on their home cage chow. The remarkable result was that there was no difference in lever press rates in these two groups. However, rats that had earlier been exposed to the pellets (in the absence of any opportunity to lever press) while they were in the sated state ultimately decreased their lever press rate during the test session when tested sated compared to when tested hungry. The result shows that in order for a reduction in drive to affect instrumental responding the animals must first learn to update the value that they attach to the food reward in that reduced drive state. The inferential reasoning account works here if it is assumed that the "desire" is what is targeted by the incentive learning manipulation. In other words, the rats are assumed to acquire the belief that "lever press leads to pellets" and they normally have the desire to "obtain pellets because they are valuable." Following a reduction in drive, this analysis does not assume that the "desire" has been altered in any way, unless the rats are given explicit training to the effect that "pellets are not valuable in the reduced drive state." Therefore, only rats that have received the relevant incentive learning treatment will show sensitivity to shifts in drive.

It may be noted that this incentive learning effect might also be described in associative terms (see Dickinson & Balleine, 1994). Indeed, whether the "belief-desire" propositional framework advanced by Mackintosh and Dickinson (1979) can be understood at the lower level of an associative theory is an interesting topic in its own right. It is not my aim to provide such an analysis. For the present purposes, I merely wish to distinguish these two types of explanations as being cast at different levels of analysis. Ultimately, as I have suggested earlier, it would be most useful if working at one level of analysis (e.g., inferential reasoning) could be reduced to another (e.g., the associative).

These two sorts of explanation, though, seem worlds apart. On the one hand, the basic mechanisms are assumed to reside in the nature of the associative structure thought to underlie learning. In principle, it should be possible for us to look in the nervous system and see just where these associations may reside. In

contrast, the basic mechanisms assumed in the inferential reasoning account are propositional encodings of associations. It seems quite a bit more difficult to imagine how an analysis of neural circuits will ultimately provide us with evidence relevant to this claim. Nevertheless, the approach has led to new and different ways of considering animal behavior (see also Blaisdell et al., 2006). For the present, I would like to ask whether evidence suggests that rats do in fact behave in essentially rational ways. There are at least two situations that would appear to present problems for this sort of approach.

Both examples of non-rational behavior come from studies of Pavlovian-to-instrumental transfer. The first problem arises from the consideration that sometimes Pavlovian stimuli exert outcome-general effects on performance. When a stimulus evokes a specific representation of an outcome and this, in turn, selectively guides instrumental response choice (see Corbit & Balleine, 2005 above), this selective transfer has a ready explanation in terms of either the associative or inferential reasoning approaches. However, the problem arises when a stimulus evokes a representation of an outcome that is unrelated to the outcomes associated with either response. As Corbit and Balleine (2005) demonstrated, such a cue generally activates all instrumental responses. But why should this be the case if decisions to respond or not are rational ones? Consider an earlier study by Dickinson and Dawson (1987). In this study rats were initially trained to press a lever for food pellets while hungry. They also learned to associate CS1 with pellets and CS2 with liquid sucrose. Subsequently, rats were given a Pavlovian-to-instrumental transfer test while either hungry or thirsty. Rats tested hungry displayed somewhat more lever pressing when tested with CS1 than with CS2. This makes sense because CS1 was previously paired with pellets as was the lever press response. However, rats tested thirsty displayed more lever press responses in the presence of CS2 compared to CS1. This does not make sense from an inferential reasoning perspective because a signal for liquid sucrose is being shown to elevate responding, ostensibly, for dry pellets while the rats are thirsty.

The associative approach, however, faces no difficulty in explaining this finding. Assume that CS2, the cue paired with sucrose, activates a representation of liquid sucrose that is motivationally more activating while the animals are thirsty than would CS1, the cue paired with pellets. This being the case, the sucrose-paired cue could merely energize responding more effectively than the pellet-paired cue.

The second instance of seemingly non-rational behavior comes from studies that have looked at the possibility that outcome-selective instrumental transfer can be shown to occur even when the outcome upon which that transfer is based itself has been devalued. One study was reported by Rescorla (1994). In this study, two instrumental responses were each trained independently with distinct reinforcing outcomes (R1-O1, R2-O2). In addition, a third instrumental response was reinforced with each of these outcomes, but which outcome that response earned depended upon which of two discriminative stimuli was present (S1: R3-O1, S2: R3-O2). The first two instrumental responses were then trained with a third outcome (R1-O3, R2-O3), before one of the outcomes was then devalued with LiCl pairings (O1-LiCl, O2-). Finally, R1 and R2 were tested in the presence of S1 and S2. In this test, R2 was responded to more in the presence of S2 than S1.

This makes perfect sense: R2 had earlier been trained with O2, and S2 had earlier signaled reward by O2. Furthermore, O2 is valuable at the time of the test. However, what was surprising was that R1 was responded to more in the presence of S1 than S2 to the same degree. This is surprising because although both S1 and R1 were previously reinforced with O1, this outcome is now devalued. From an inferential reasoning perspective it makes no sense why the animal should respond more in the presence of a cue with the hopes of obtaining an outcome that is devalued. On the other hand, the associative perspective would not have particular trouble with this finding if it were assumed that the bidirectional R-O link operated through the specific sensory properties of the outcome (rather than its value).

To be sure there is much we do not know concerning the structures that control instrumental behavior. However, these two examples present special challenges to an inferential reasoning account. Of course, this does not invalidate the inferential reasoning approach to instrumental behavior, but it does add constraints to the assumptions that would need to be made in order for such an account to explain the full range of phenomena. It seems possible, for instance, that one might entertain an inferential reasoning explanation that could make sense of the two findings noted above, particularly if the basic premises of the propositional syllogism were encoded in a way that did not fully capture the richness of the environmental contingencies. For instance, if the rat learns that "pressing the lever leads to something good" rather than "pressing leads to a pellet" then when the rats are tested in a state of thirst a stimulus that signals liquid sucrose, itself, could be encoded as "this stimulus signals that something good can now be obtained." Putting together these two premises could lead the rat to seemingly irrationally press the pellet-reinforced lever more in the presence of the cue for sucrose. However, caution should be applied here because by assuming such degraded propositional encodings, this sort of analysis will run into trouble with other quite basic effects, such as reward-specific Pavlovian-to-instrumental transfer effects. Nevertheless, the main point is that an inferential reasoning approach that assumes rational decision-making will have difficulty explaining the two phenomena I have described above without making additional assumptions. Ultimately, the theorist will need to assess whether this type of explanation works best by also considering the full range of assumptions that would need to be made to capture the phenomena of interest.

The associative approach, as well, faces significant challenges. Without going into any great detail regarding the associative structures thought to underlie the control of instrumental performance, I present a simple diagram (Fig. 2) that illustrates the main challenge for an associative theory of instrumental learning to solve. In this diagram, I assume that the brain will in some manner represent the stimulus situation, the response, and the outcome (S, R, O, respectively). There are a variety of potential associative connections that may form between these three terms. Indeed, throughout the history of the study of instrumental learning theorists have envisioned many different solutions. These include the simple S-R association of Thorndike mentioned above, the R-O association initially emphasized by Tolman, the S-O, R-O bidirectional chain model of Pavlov, plus several others that I have not mentioned. One of these was introduced by Rescorla and Solomon (1967), the S-O, S-R energization model, another by Trapold and

Overmier (1972), the S-O, O-R chain model, and the S-[R-O] hierarchical learning model, introduced by Colwill and Rescorla (1990). There exists reasonable support for each of these associative accounts, and it is not the place here for me to review these findings. Instead, I would like to point out that one major challenge as I see it is in coming up with a unified approach for understanding how all of these associative processes may interact with one another.

There is a current line of research that is beginning to actively pursue at least part of this problem. It concerns the issue of whether instrumental responding undergoes a transition from being "goal-directed" early in training to "habitual" after extended training (e.g., Adams & Dickinson, 1981). By these terms it is usually meant that the response undergoes a transition from being controlled by R-O associations early to S-R associations late in training. The main challenge here is to understand how this works, and to describe clearly the nature of the system that gives rise to it. Recent studies conducted at the neural systems level of analysis have revealed interesting facts about the process (e.g., Killcross & Coutureau, 2003), but, nevertheless, some puzzles remain. One particularly intriguing problem is in understanding why under some conditions goal-directed responding never seems to become habitual (e.g., Colwill & Triola, 2002).
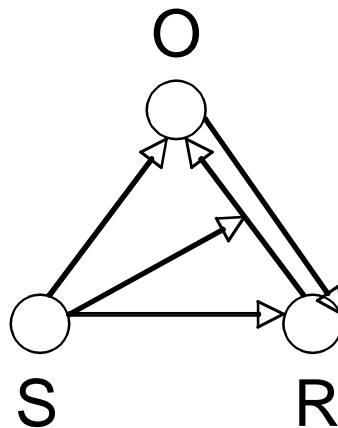
Figure 2. This figure displays the different possible relationships among stimulus (S), response (R), and reinforcing outcome (O) in instrumental conditioning. See text for a further description of these relationships.

Some Thoughts on Causal Reasoning in Rats

In an intriguing set of studies it has been claimed that rats possess the ability to display causal reasoning (e.g., Blaisdell et al., 2006; Leising, Wong, Waldmann, & Blaisdell, 2008). In the basic experimental design the rats are first given pairings of two neutral stimuli, e.g., a light stimulus is followed by a tone stimulus (L – T). Subsequently, rats are given first order conditioning with the first of these stimuli, i.e., they are given L – sucrose pairings. In essence, the rats could potentially learn that both T and sucrose events are effects of a common cause (L). During a test phase, if rats have this "causal model" then when T is presented they would attribute its appearance to the presence of L, which, in turn, should lead to

the expectation that the other effect (sucrose) should also occur. However, if during the test phase an instrumental contingency is imposed of the form that whenever the rats press a lever this response produces T, then rats have reason to believe that T was produced by this response and not by L. This would have the effect of breaking the link between T and sucrose as having a common cause. Thus, rats should display more magazine responses during T when it was presented alone (the "observe" condition) than when it was presented following the lever press response (the "intervene" condition). This is exactly the result that the authors obtained in several studies (but see Dwyer, Starns, & Honey, 2009).

The authors suggested that this result cannot easily be understood from the perspective of an associative theory partly on the basis of the additional finding that this "intervene" manipulation had no effect on responding to T when it had originally been trained as part of a causal chain (i.e., T – L, then L – Sucrose, then test T versus lever press – T). The comparison between these two conditions is an interesting one because very similar procedures are used in the two cases with the exception that the order in which L and T are presented during their training trials differs (L – T versus T – L). However, this subtle difference in procedure, apart from changing the model from a "common cause" model to a "causal chain" model, also has consequences from the point of view of an associative theory.

Consider what might happen when one trains with the "common cause" procedure. First, L – T trials establishes an association between the two events. Subsequently, when L – sucrose trials are introduced, then the prior L – T association could enable L to associatively activate a representation of T at the time it is paired with sucrose. This could result in "mediated conditioning" occurring to T, because its representation (along with L) is actually paired with sucrose. Thus, responding to this stimulus during the test session could reflect this type of learning. Holland (1998) demonstrated that mediated conditioning of this sort is possible. However, the situation is different in the "causal chain" scenario. Here, one might expect responding during the test to reflect, as the authors suggest, a causal chain of associations: T – L, L – sucrose. The significance of these considerations is that the analysis points to different mechanisms responsible for producing responding to T during the test session in the two situations. However, in order to explain the effects of the intervention manipulation an additional consideration needs to be made. To be concrete, why might a lever press response that occurs before T diminish responding to T? One possibility is that this lever press response requires some degree of processing within a limited capacity processing system (e.g., Wagner, 1981). If this processing interferes with processing of the T when it is presented immediately after the response is made, then there will be less effective processing of T and this will reduce its ability to associatively activate its associated sucrose representation, thereby, reducing magazine responding in the "intervene" condition. Now, for this analysis to account for the findings, one would need to additionally assume that mediated conditioning is more subject to processing interference than is associative chaining. Given that we do not understand well the various properties of mediated learning and associative chaining processes, it would seem hard to rule out this interpretation of the findings at the present time.

One potential way to test the differences between these two approaches might be to simply increase the total amount of training. It is hard to see why more extensive training should have any impact if a causal model framework is adopted because the same common cause, multiple effect model would apply after limited or extended amounts of training. However, according to an associative theory, one would expect that after many interspersed L-sucrose, L-T trials, the T should acquire conditioned inhibitory properties because T would eventually lead the subject to expect the absence of sucrose on that trial. In other words, the mediated conditioning process noted above should eventually give way to a more potent conditioned inhibition process. Such effects have been observed in closely related second order conditioning situations (e.g., Stout, Escobar, & Miller, 2004). The upshot of this is that an associative model might very well apply to the intriguing findings noted above.

Conclusions

In this article I have tried to illustrate the richness of an associative learning perspective in accounting for a wide variety of research results that can be characterized as involving "cognition." I hope we are well past the time when associative theories will be dismissed out of hand because of a historical predilection against such analyses. The successful application of connectionist theories to human cognition together with an increasing understanding at the neural level of analysis of how associative systems work both point to the fruitfulness of the approach. Indeed, purely behavioral studies have revealed some interesting facts about associative learning and some of its cognitive underpinnings. At the heart of the approach is the notion that representations of events enter into the association. There is much current work directed at understanding the nature of reward representations in Pavlovian learning and some of this work is revealing (1) how "images" and "expectancies" may differ in their control of behavior, (2) how the representation of "what" US to expect and "when" it might appear may involve different learning systems, and (3) that learning "what" US to expect and about its motivational value also seems to involve separate systems. These distinct forms of US representation (image, expectancy, temporal, motivational) reflect just four aspects of the complex events we refer to as USs. Other components of the US include their hedonic qualities and response evoking potentials. Presumably associations may also be formed with these other aspects of the US, however, we currently lack very clear ways of assessing their existence. Future research will profit by developing ways of identifying them.

The study of instrumental learning has also led to a very rich way of thinking about behavior from the associative perspective. Although there are only a few terms in the equation (S, R, O) as the above discussion has made clear there are a wide variety of ways of thinking about how those three terms might coalesce within various associative structures. Undoubtedly, there are multiple associative structures that underlie instrumental learning and performance. One major challenge will be to understand in a more unified manner how these various associative structures participate together to influence behavior. Studies examining the transition from control by R-O to S-R associations (i.e., goal-directed to

habitual behavior) are making headway into this problem, but the problem will also need to consider some of the other associative structures noted above.

I have also considered the potential role of an inferential reasoning account of learning. This approach is based on assumptions quite different from the associative approach, however, to the extent that unique predictions can be derived from the approach it will undoubtedly stimulate further work. There are places where an inferential reasoning perspective will likely fall short in accounting for certain aspects of instrumental performance, namely, with outcome-general transfer and Pavlovian-instrumental transfer mediated by a devalued outcome. Nevertheless, as the recent work reviewed above of Blaisdell and his colleagues has shown the approach can lead one to perform interesting experiments that might not have been entertained through an associative perspective. At the same time I have indicated how one associative account may be applied to the findings that gave rise to the claim of "causal reasoning" in rats in the first place. Whether or not this analysis will be useful will depend upon further research.

Finally, in discussing the differences between "learning" on the one hand and "cognition" on the other, I will have had some impact if I have made clear that this boundary just might turn out in the end to be a rather imaginary one. The associative tradition in the study of mind has a long history and has been extremely successful, in my view, in providing a realistic perspective on how basic cognitive and brain functions might operate. Future research will undoubtedly reveal many new discoveries about how this remarkable learning system works and how rather complex behavior might emerge from a much simpler set of principles.

References

Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *Quarterly Journal of Experimental Psychology B*, *33*, 109-121.

Atkinson, R. C., & Shiffrin, R. M. (1971). The control of short-term memory. *Scientific American*, *225*, 82-90.

Balleine, B. (1992). Instrumental performance following a shift in primary motivation depends on incentive learning. *Journal of Experimental Psychology:Animal Behavior Processes*, *18*, 236-250.

Balleine, B. & Ostlund, S. B. (2007). Still at the choice-point: Action selection and initiation in instrumental conditioning. *Annual Review of the New York Academy of Sciences*, *1104*, 147-171.

Balsam, P. D. & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neuroscience*, *32*, 73-78.

Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, *311*, 1020-1022.

Boakes, R. (1984). *From Darwin to behaviorism: Psychology and the minds of animals*. Cambridge, UK: Cambridge University Press.

Bouton, M. E., & Garcıa-Gutierrez, A. (2006). Intertrial interval as a contextual stimulus. *Behavioural Processes, 71,* 307-317.

Bouton, M. E., & Hendrix, M. C. (2011). Intertrial interval as a contextual stimulus: Further analysis of a novel asymmetry in temporal discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *37*, 79-93.

Corbit, L. H., & Balleine, B. W. (2005). Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of Pavlovian-instrumental transfer. *Journal of Neuroscience*, *25*, 962-970.

Colwill, R. M., & Motzkin, D. K. (1994). Encoding of the unconditioned stimulus in Pavlovian conditioning. *Animal Learning & Behavior*, *22*, 384-394.

Colwill, R. M., & Rescorla, R. A. (1985). Postconditioning devaluation of a reinforcer affects instrumental responding. *Journal of Experimental Psychology:Animal Behavior Processes, 11*, 120-132.

Colwill, R. M., & Rescorla, R. A. (1990). Evidence for the hierarchical structure of instrumental learning. *Animal Learning & Behavior*, *18*, 71-82.

Colwill, R. M., & Triola, S. M. (2002). Instrumental responding remains under the control of the consequent outcome after extended training. *Behavioral Processes*, *57*, 51-64.

Delamater, A. R. (1998). Associative mediational processes in the acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes, 24*, 467-482.

Delamater, A. R., & Holland, P. C. (2008). The influence of CS-US interval on several different indices of learning in appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *34*, 202-222.

Delamater, A. R., & LoLordo, V. M. (1991). Event revaluation procedures and associative structures in Pavlovian conditioning. In L. Dachowski & C. F. Flaherty (Eds.), *Current topics in animal learning: Brain, emotion, and cognition* (pp. 55-94). Hillsdale, NJ: Lawrence Erlbaum Associates.

Delamater, A. R., LoLordo, V. M., & Berridge, K. C. (1986). Control of fluid palatability by exteroceptive Pavlovian signals. *Journal of Experimental Psychology Animal Behavior Processes, 12*, 143-152.

Delamater, A. R., & Oakeshott, S. (2007). Learning about multiple attributes of reward in Pavlovian conditioning. *Annual Review of the New York Academy of Sciences, 1104,* 1-20.

Delamater, A. R., Sosa, W., & Katz, M. (1999). Elemental and configural processes in patterning discrimination learning. *Quarterly Journal of Experimental Psychology B, 52*, 97-124.

Desgranges, B., Ramirez-Amaya, V., Ricaño-Cornejo, I., Levy, F., & Ferreira, G. (2010). Flavor preference learning increases olfactory and gustatory convergence onto single neurons in the basolateral amygdala but not in the insular cortex in rats. *Plos One, 5,* 1-8.

Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge, UK: Cambridge University Press.

Dickinson, A. (1989). Expectancy theory in animal conditioning. In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 279-308). Hillsdale, NJ: Lawrence Erlbaum Associates.

Dickinson, A., & Balleine, B. W. (1994). Motivational control of goal-directed action. *Animal Learning & Behavior*, *22*, 1-18.

Dickinson, A., & Dawson, G. R. (1987). Pavlovian processes in the motivational control of instrumental performance. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 39B*, 201-213.

Dickinson, A., & Dawson, G. R. (1988). Motivational control of instrumental performance: The role of prior experience of the reinforcer. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology, 40B,* 113-134.

Dwyer, D. M., Starns, J., & Honey, R. C. (2009). "Causal reasoning" in rats: A reappraisal. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 578-586.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, *107*, 289-344.

Guthrie, E. R. (1935). *The psychology of learning*. New York: Harper.

Hall, G. (1996). Learning about associatively activated stimulus representations: Implications for acquired equivalence and perceptual learning. *Animal Learning & Behavior, 24*, 233-255.

Hall, G. (2002). Associative structures in Pavlovian and instrumental conditioning. In R. Gallistel (Ed.), *Stevens' handbook of experimental psychology, Volume 3: Learning, motivation, and emotion* (pp. 1-45). New York: John Wiley & Sons, Inc.

Holland, P. (1998). Amount of training affects associatively-activated event representation. *Neuropharmacology, 37*, 461-469.

Holland, P. (1990). Event representations in Pavlovian conditioning: Image and action. *Cognition, 37*, 105-131.

Holland, P. C., Lasseter, H., & Agarwal, I. (2008). Amount of training and cue-evoked taste-reactivity responding in reinforcer devaluation. *Journal of Experimental Psychology: Animal Behavior Processes, 34*, 119-132.

Kehoe, E. J., Horne, P. S., Macrae, M., & Horne, A. J. (1993). Real-time processing of serial stimuli in classical conditioning of the rabbit's nictitating membrane response. *Journal of Experimental Psychology: Animal Behavior Processes, 19,* 265-283.

Kerfoot, E. C., Agarwal, I., Lee, H. J., & Holland, P. C. (2007). Control of appetitive and aversive taste-reactivity responses by an auditory conditioned stimulus in a devaluation task: A FOS and behavioral analysis. *Learning and Memory, 14*, 581-589.

Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral Cortex, 13*, 400-408.

Konorski, J. (1967). *Integrative activity of the brain*. Chicago: University of Chicago Press.

Leising, K. J., Wong, J., Waldmann, M. R., & Blaisdell, A. P. (2008). The special status of actions in causal reasoning in rats. *Journal of Experimental Psychology: General, 137,* 514-527.

Mackintosh, N. J. (1974). *The psychology of animal learning*. London: Academic Press Inc.

Mackintosh, N. J. (1983). *Conditioning and associative learning*. New York: Oxford University Press.

Mackintosh, N. J., & Dickinson, A. (1979). Instrumental (Type II) conditioning. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation* (pp. 143-167). Hillsdale, NJ: Erlbaum.

Meck, W. H., & Church, R. M. (1982). Abstraction of temporal attributes. *Journal of Experimental Psychology: Animal Behavior Processes, 8*, 226-243.

Nadler, N., Delgado, M. R., & Delamater, A. R. (2011). Pavlovian to instrumental transfer of control in a human learning task. *Emotion, 11*, 1112-1123.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.

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

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.

Randich, A., & LoLordo, V. M. (1979). Associative and non-associative theories of the UCS preexposure phenomenon: Implications for Pavlovian conditioning. *Psychological Bulletin, 86*, 523-548.

Rescorla, R. A. (1994). Transfer of instrumental control mediated by a devalued outcome. *Animal Learning and Behavior, 22*, 27-33.

Rescorla R. A., & Solomon, R. L. (1967). Two-process learning theory: Relationship between Pavlovian conditioning and instrumental learning. *Psychological Review*, *74*, 151-182.

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

Reynolds, S. M., & Berridge, K. C. (2002). Positive and negative motivation in nucleus accumbens shell: Bivalent rostrocaudal gradients for GABA-elicited eating, taste "liking"/"disliking" reactions, place preference/avoidance, and fear. *Journal of Neuroscience*, *22*, 7308-7320.

Rumelhart, D. E., McClelland, J. L., & the PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: The MIT Press.

Saddoris, M. P., Holland, P. C., & Gallagher, M. (2009). Associatively learned representations of taste outcomes activate taste-encoding neural ensembles in gustatory cortex. *The Journal of Neuroscience, 29,* 15386-15396.

Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.

Stout, S., Escobar, M., & Miller, R. R. (2004). Trial number and compound stimuli temporal relationship as joint determinants of second-order conditioning and conditioned inhibition. *Learning & Behavior*, *32*, 230-239.

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. New York: Macmillan.

Tolman, E. C. (1932). *Purposive behavior in animals and men*. New York: Appleton-Century-Crofts.

Trapold, M. A., & Overmier, J. B. (1972). The second learning process in instrumental learning. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 427-452). New York: Appleton-Century-Crofts.

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.

Wagner, A. R., & Brandon, S. E. (1989). Evolution of a structured connectionist model of Pavlovian conditioning (AESOP). In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 149-189). Hillsdale, NJ: Erlbaum.