Special Section

Computational Versus Associative Models of Simple Conditioning

C.R. Gallistel[1] and John Gibbon[2]

Department of Psychology and the Center for Cognitive Science, Rutgers University, New Brunswick, New Jersey (C.R.G.), and Division of Biopsychology, New York State Psychiatric Institute, New York, New York, and Department of Psychology, Columbia University, New York, New York (J.G.)

Abstract

In associative models of simple conditioning, conductive connections (associations, Hebbian synapses) are strengthened by the repetitive temporal pairing of stimuli. The associations cause the animal to behave more adaptively, but they do not encode information about objectively specifiable properties of the conditioning experience. In information processing (computational) models, the temporal intervals in that experience are timed and the results recorded in memory for later use in computations that determine the decisions whether and when to respond to the conditioned stimulus. The predictions of these latter models depend on the ratios of remembered and currently experienced temporal intervals; hence, they are time-scale invariant. Two examples of empirical time-scale invariance are described: Neither the delay of reinforcement nor the ratio of reinforced to unreinforced presentations of the conditioned stimulus affects rates of acquisition and extinction. Time-scale invariance has far-reaching implications for models of the processes that underlie conditioning, for example, models of Hebbian synapses.

Keywords

learning; associations; information processing; time-scale invariance; Hebbian synapses

The concept of an associative bond has been central to psychologists' understanding of learning for more than a century. Computational theories of mind—also called information processing theories—offer a fundamentally different way to think about learning, a way in which this concept plays no role. The differences between the frameworks are clearly seen in the contrasting accounts they offer for Pavlovian conditioning, an experimental paradigm created in order to determine the laws of association formation. Experiments using this paradigm have suggested a quantitative property of the conditioning process—time-scale invariance—that is deeply difficult to reconcile with the fundamental assumptions of associative theory but follows directly from the fundamental assumptions of our information processing models.[3]

In Pavlovian conditioning, a behaviorally neutral stimulus, called the conditioned stimulus (CS), is repeatedly paired with a motivating stimulus, called the unconditioned stimulus (US), until the subject responds to the CS in a manner that anticipates the US. For example, in the method most often used to condition a hungry pigeon to peck a key for food, the CS is the illumination of a round key, which remains illuminated for a few seconds (the CS-US interval), until a food hopper opens for a few seconds (the US). This sequence of events is repeated after some interval (the intertrial interval) until the pigeon begins to peck at the illuminated key. Note that the food appears at a certain latency (the delay of reinforcement) whether the pigeon pecks the key or not, making this a Pavlovian procedure, not an instrumental procedure.

CONTRASTING EXPLANATIONS

In associative models (e.g., Rescorla & Wagner, 1972), the temporal pairing of the key's illumination and food creates new conductive links, called associations. Neurobiologists suppose that these new connections are modified synapses between neurons, called Hebbian synapses. Conductance at these synapses is modified by the temporal pairing of pre- and postsynaptic activity, with the CS causing the presynaptic activity and the US the postsynaptic activity. Association formation is progressive; successive temporal pairings of the CS and the US strengthen (reinforce) the same association. Thus, the current strength of an association reflects many different aspects of the animal's conditioning experience, which means that it does not represent any objective aspect of that experience, such as the CS-US interval. An association is not a symbol; its strength does not represent any objective fact about the world, and the associative bond does not participate in information processing (computational) operations. Associations, unlike symbols, are not added, subtracted, multiplied, and divided in order to generate new associations.

In the contrasting information processing account that we consider here (Gallistel & Gibbon, 2000), the pigeon's brain times the durations of temporal intervals, such as the CS-US interval, and records the results in memory for later use in the computations that mediate decisions on whether to respond to a CS and how long to delay that response.

In the rate-estimation theory (RET) of acquisition, conditioned responding appears only when the pigeon has decided that there is a strong enough CS-US contingency. Its measure of contingency is the ratio of its estimates for two interreinforcement intervals, the expected interval between feedings when the key is illuminated (CS present) and the expected interval between feedings when it is not (CS absent). During the intertrial intervals, when the CS is absent, the only stimulus is the experimental chamber itself, which is the background or context in which conditioning occurs. Thus, the pigeon's measure of contingency is the ratio between its estimate of the background interreinforcement interval and the interreinforcement interval when the CS is present. Its estimate of the latter interval is determined by T, which is the interval between the onset of the CS and the delivery of the food, at CS offset. This interval is commonly called the delay of reinforcement or the trial duration. The expected interval between reinforcements when the CS is present is T times the expected number of trials (CS presentations) per reinforcement. This latter number (actually a ratio) is S, the partial reinforcement schedule.

In the scalar-expectancy theory (SET) model of response timing, a pigeon that has already decided that there is a strong enough CS-US contingency nonetheless refrains from pecking at the illuminated key until a certain proportion of T has elapsed. That is, the timing of the onset of the conditioned response relative to the onset of the CS is controlled by the ratio between the remembered delay of reinforcement and the currently elapsed interval since the onset of the CS.

In these information processing accounts, learning is the process of computing from raw sensory input objective properties of the experienced world and storing the results in a memory. The memories thus created are not conducting links; they are repositories of information, like the bit patterns in the memory of a conventional computer or the genes on a chromosome. Like computer memories and genes, they must be read in order to have an influence on observable output. Unlike associations, memories specify objective properties of the animal's experience. What is stored in memory is the information extracted from experience. Also unlike associations, these memories enter into computational operations.

The radically different nature of the two accounts may be appreciated by considering the answers they offer to the kinds of questions commonly addressed in textbooks on animal learning (see Table 1).

TIME-SCALE INVARIANCE

The value of reconceptualizing a long-familiar phenomenon lies in the extent to which the new conceptual framework leads to experiments that deepen insight into the fundamental nature of the phenomenon. The information processing framework brings into sharp focus a quantitative principle about conditioning that, if generally true, is profoundly important. The principle of time-scale invariance asserts that the time scale of an experimental protocol—the absolute durations of the temporal intervals that define the protocol—does not affect the outcome of the experiment, because only the proportions (ratios) among the intervals in the protocol matter. Elsewhere (Gallistel & Gibbon, 2000), we have discussed many manifestations of this principle. Here we describe two that have strong implications for associative models, particularly those that describe the association-forming process in physiological terms, by specifying a physiological mechanism, like long-term potentiation, for realizing a Hebbian synapse (e.g., Brown, Kairiss, & Keenan, 1990; Magee & Johnston, 1997).[4]

One manifestation of the time-scale invariance of the conditioning process is that the delay of reinforcement has no effect on the rate of conditioning, over a wide range of delays—provided that the other intervals in the experimental protocol are varied in proportion to the variation in the delay of reinforcement (Gibbon, Baldock, Locurto, Gold, & Terrace, 1977; see Fig. 1). Another manifestation is the lack of an effect of partial reinforcement—either on the number of reinforcements required for the acquisition of a conditioned response or the number of reinforcements that must subsequently be omitted in order to extinguish (eliminate) the conditioned response (Gibbon, Farrell, Locurto, Duncan, & Terrace, 1980; see Fig. 2). In a partial-reinforcement protocol, reinforced trials and unreinforced trials—trials on which the CS is presented but the US is omitted—are randomly intermingled during training. The greater the fraction of unreinforced trials, the thinner the schedule of reinforcement, S.

In associative models, nonreinforcement of a CS weakens the effects of reinforcement. Thus, interpolating many unreinforced trials during training should increase the number of reinforced trials required for acquisition. In fact, however, partial reinforcement has little or no effect, even when there are, on average, as many as nine

Table 1. *Different answers to basic questions*

• Why does the conditioned response appear?
AM: The response appears because the associative connection gets stronger.
CM: The response appears because the ratio of the estimated rate of CS reinforcement to the estimated rate of background reinforcement grows until it exceeds the decision threshold.

• Why does the conditioned response disappear during extinction, when the CS no longer predicts the US?
AM: The response disappears because there is a loss of net excitatory associative strength.
CM: The response disappears because the ratio between the interval elapsed since the last CS reinforcement and the expected interval between CS reinforcements grows until it exceeds a decision threshold.

• What is the effect of reinforcement (US delivery)?
AM: It strengthens excitatory associations.
CM: It starts or stops one or more timers.

• What is the effect of delay of reinforcement?
AM: It reduces the increment in associative strength produced by a reinforcement.
CM: It lengthens the remembered interreinforcement interval in the presence of the CS, the remembered CS-US interval, or both.

• What is the effect of nonreinforcement (withholding the US, a physically undefinable event)?
AM: This physically undefinable event somehow weakens an excitatory association, or it strengthens an inhibitory association; in either case, it reduces the net excitatory effect of the CS.
CM: This nonevent has no effect; the timer for the most recent interreinforcement interval continues to accumulate, because nothing has happened.

• What is the effect of varying the magnitude of reinforcement (amount of food)?
AM: It varies the size of the increment in the excitatory association.
CM: It varies the remembered magnitude of reinforcement.

• What happens when nothing happens (during an intertrial interval)?
AM: Nothing.
CM: The timer that times the duration of the animal's experience of the experimental chamber continues to accumulate, steadily decreasing the animal's estimate of the background rate of reinforcement (the rate in the absence of the CS).

• What is the effect of the onset of the CS?
AM: It opens the associative window in the mechanism that responds to the temporal pairing of two signals (the Hebbian synapse); that is, it begins a trial during which the updating of associative strength will occur.
CM: It starts a timer (to time the duration of this presentation), and it causes the cumulative CS-exposure timers to resume cumulating.

• What happens when more than one CS is present during reinforcement?
AM: The CSs compete for a share of a limited increment in associative strength, or selective attention to one CS denies other CSs access to the associative mechanism.
CM: The rate of reinforcement is partitioned among reinforced CSs in accord with two computational principles—additivity and predictor minimization—to yield rate-of-reinforcement estimates for each CS.

• How does conditioned inhibition[a] arise?
AM: The omission of an otherwise expected US (the occurrence of what is sometimes called a no-US event) strengthens an inhibitory association.
CM: The animal estimates that the CS reduces the rate of reinforcement, an effect equivalent to adding a negative rate of reinforcement to whatever positive rate is otherwise in force.

• Why is the delay between the onset of the CS and appearance of a conditioned response proportional to the delay between the onset of the CS and the onset of the US?
AM: There is no widely accepted answer to this question in associative theory.
CM: The two delays are proportional because the animal remembers the delay of the US, and it compares the currently elapsing interval since the onset of the CS to that remembered delay and responds only when the currently elapsed interval approximates the remembered delay.

• What happens when a CS follows the US rather than preceding it (backward conditioning)?
AM: Nothing; only forward temporal pairing produces associations [according to some models]. An inhibitory connection between CS and US is formed [according to other models].
CM: A negative CS-US interval or, equivalently, a positive US-CS interval is recorded. (Remembered intervals, like remembered rates, are signed.)

(*Continued*)

Table 1. *Continued*

• How does a secondary CS[b] acquire behavioral potency?
AM: An association forms between the secondary CS and the primary CS, so that activation may be conducted from the secondary CS to the primary CS and thence to the US.
CM: The (signed) interval between the secondary and primary CSs is summed with the (signed) interval between the primary CS and the US to obtain the expected interval between the secondary CS and the US.

Note. AM = associative model; CM = computational model; CS = conditioned stimulus; US = unconditioned stimulus.
[a]Conditioned inhibition refers to the case in which the animal learns to inhibit responding to CSs that predict the withholding of the US.
[b]A secondary CS predicts a primary CS, which alone has been directly paired with the US.

unreinforced trials for every reinforced trial (i.e., when $S = 10$; Fig. 2; see also Williams, 1981). In our model of acquisition (Gallistel & Gibbon, 2000), the subject computes the expected interval between reinforcements both when the CS is present and during intertrial intervals, when it is absent. There are no reinforcements during the intertrial intervals, so in our model the estimate for this interevent interval grows in proportion to the cumulative intertrial interval, that is, in proportion to the cumulative duration of the interval of observation.[5] The subject responds to the CS when the ratio of these two estimates exceeds a decision criterion. With a partial-reinforcement schedule with S trials for every reinforced trial, it takes S times more trials to reach some critical number of reinforcements than if all trials are reinforced, but the ratio of the two estimated interreinforcement intervals is the same when that critical number is reached. Thus, the lack of an effect of partial reinforcement on reinforcements to acquisition is another manifestation of the time-scale invariance of the conditioning process.

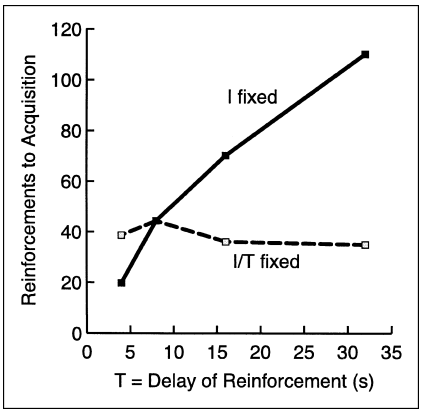

Fig. 1. Number of reinforcements (feedings) required for the onset of reliable conditioned responding in experiments by Gibbon, Baldock, Locurto, Gold, and Terrace (1977). The experiments varied both the delay of reinforcement, T (the interval between onset of the conditioned stimulus [CS] and delivery of the unconditioned stimulus [US], which coincides with the offset of the CS), and the intertrial interval, I (the interval between presentations of the CS). When I was kept constant while T was increased, thereby changing the proportions between these two intervals, the number of reinforcements required increased. (Equivalently, the rate of acquisition decreased.) When, however, I was increased in proportion to T, thereby holding fixed the I/T ratio, the increase in T had no effect on the rate of acquisition. (From Gallistel & Gibbon, 2000, by permission of the publisher.)

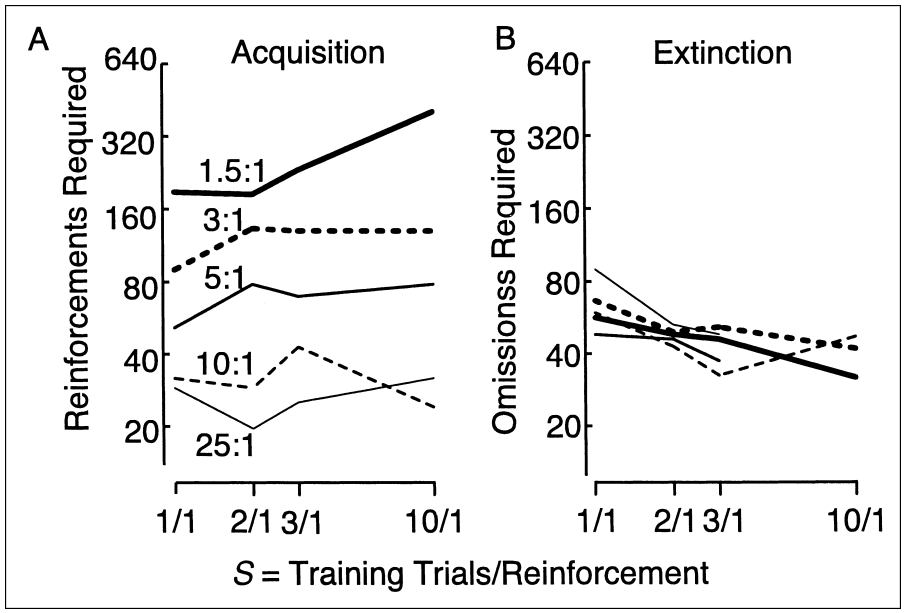

Fig. 2. Effects of the I/T ratio (ratio of the intertrial interval to the delay of reinforcement) and the reinforcement schedule on acquisition (a) and extinction (b) in the Pavlovian conditioning of pigeon key pecking. The graph in (a) shows the number of reinforcements required to establish conditioned responding. The number pair by each curve (e.g., 5:1) gives the I/T ratio. The bigger this ratio, the fewer the reinforcements required for acquisition. By contrast, the reinforcement schedule (S = trials/reinforcement) has little or no effect. The graph in (b) shows the number of expected reinforcements that must be omitted to extinguish conditioned responding as a function of the I/T ratio and S. Neither variable has much effect. (Replotted from data in Fig. 14 of Gallistel & Gibbon, 2000.)

Our computational models are naturally time-scale invariant, because they are built on the ratios of remembered or currently measured temporal quantities; changing the time scale has no effect on these temporal ratios. Associative models, by contrast, are not time-scale invariant, because associations are assumed to form only when the CS and US are temporally paired, that is, when the US follows the CS within some critical interval. (For an example of the central role this assumption plays in neurobiologically oriented associative theorizing, see Tang et al., 1999.) Many associative models also make extensive use of intrinsically decaying stimulus traces (Sutton & Barto, 1990; Wagner, 1981); that is, the trace (aftereffect) of a stimulus within the nervous system is imagined to decay over time or become in one way or another less accessible to the associative process. The intrinsic decay rates make the predictions of these models very sensitive to the time scale of the protocol. Other associative models (e.g., Rescorla & Wagner, 1972) carve the time the subject spends in the experimental apparatus into a sequence of imaginary trials. These trials must have an assumed duration, and that assumed duration makes the predictions of these models very sensitive to the time scale of a protocol (Granger & Schlimmer, 1986).

In summary, the use of a wide variety of conditioning paradigms and subject species to determine the extent to which time-scale invariance is a general property of the conditioning process is an important new direction in the study of learning. The outcome of such a program of research may revolutionize psychologists' understanding of the process. If time-scale invariance proves to be a very general property of conditioning, this will require either quite radical reformulation of associative models or their abandonment in favor of information processing models. The implications for neurobiology are also potentially far-reaching. Long-term potentiation and long-term depotentiation, which many neurobiologists believe to be the cellular processes that mediate learning and memory (see, e.g., Tang et al., 1999), are not time-scale invariant processes.

Recommended Reading

Gallistel, C.R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.

Gallistel, C.R., & Gibbon, J. (2000). (See References)

Rescorla, R.A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist, 43*, 151–160.

Acknowledgments—The authors gratefully acknowledge the support of the following grants: SBR-9720410 from the National Science Foundation to Rochel Gelman, C.R. Gallistel, Orville Chapman, et al., co-principal investigators, and MH41649 from the National Institute of Mental Health to John Gibbon.

Notes

1. Address correspondence to C.R. Gallistel, Rutgers Center for Cognitive Science, Rutgers University, 152 Frelinghuysen Rd., Piscataway, NJ 08854-8020; e-mail: galliste@ruccs.rutgers.edu.

2. John Gibbon died after we had completed preparation of this article. I miss him sorely, both as engaging friend and as stimulating collaborator. His pioneering theoretical and experimental work on the psychophysical characteristics of memory for temporal durations placed him in the first rank among students of learning and memory.

3. Readers who would like to learn more about learning from an information processing perspective may want to consult the first Recommended Reading, which is a broad survey of the topic. The second Recommended Reading focuses on the traditional classical and instrumental conditioning literature; the third Recommended Reading emphasizes the many important discoveries about classical and instrumental conditioning that have not found their way into the textbooks.

4. Long-term potentiation and depotentiation are long-lasting changes in synaptic conductance produced by combining activation of the input to a synapse with some kind of activation or inactivation of the postsynaptic membrane.

5. The principle involved here is that when there has been no event within the interval of observation, the cumulative interval of observation is a lower limit on what the estimate of the interevent interval may be.

References

Brown, T.H., Kairiss, E.W., & Keenan, C.L. (1990). Hebbian synapses: Biophysical mechanisms and algorithms. *Annual Review of Neuroscience, 13*, 475–511.

Gallistel, C.R., & Gibbon, J. (2000). Time, rate and conditioning. *Psychological Review, 107*, 219–275.

Gibbon, J., Baldock, M.D., Locurto, C.M., Gold, L., & Terrace, H.S. (1977). Trial and intertrial durations in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes, 3*, 264–284.

Gibbon, J., Farrell, L., Locurto, C.M., Duncan, H.J., & Terrace, H.S. (1980). Partial reinforcement in autoshaping with pigeons. *Animal Learning and Behavior, 8*, 45–59.

Granger, R.H.J., & Schlimmer, J.C. (1986). The computation of contingency in classical conditioning. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 20, pp. 137–192). New York: Academic Press.

Magee, J.C., & Johnston, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science, 275*, 209–213.

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.

Sutton, R.S., & Barto, A.G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). Cambridge, MA: Bradford Books/MIT Press.

Tang, Y.-P., Shimizu, E., Dube, G.R., Rampon, C., Kerchner, G.A., Zhuo, M., Liu, G., & Tsien, J.Z. (1999). Genetic enhancement of learning and memory in mice. *Nature, 401*, 63–69.

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.

Williams, B.A. (1981). Invariance in reinforcements to acquisition, with implications for the theory of inhibition. *Behaviour Analysis Letters, 1*, 73–80.