In press Current Directions in Psychological Science

Computational Versus Associative Models of Simple Conditioning[i]

C. R. Gallistel

University of California, Los Angeles

John Gibbon

New York State Psychiatric Institute and Columbia University

Abstract

In associative models, 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 that define the protocol are timed and the results recorded in memory for later use in the computation of the decision variables on which conditioned responding is based. 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 CS presentations appear to affect 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.

Learning Associations Information Processing Time Scale Invariance

Hebbian synapses

The concept of an associative bond has been central to our 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, 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.

---

[i] [1]Address correspondence to C.R. Gallistel, Department of Psychology, UCLA, Los Angeles, CA 90095-1563. The authors gratefully acknowledge the support of the following grants: SBR-9720410 from the NSF to Rochel Gelman, C. R. Gallistel, Orville Chapman, et al, CoPIs and MH41649 from NIMH to John Gibbon

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 respond 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.

Contrasting Explanations

In associative models (e.g., Rescorla & Wagner, 1972), the temporal pairing of the key 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 post-synaptic activity. Association formation is progressive; successive temporal parings 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, for example, the CS-US interval. An association is not a symbol; its strength does not encode information about the conditioning experience, and the associative bond does not participate in information processing (computational) operations. Associations, unlike symbols, are not added, subtracted, mutiplied and divided in order to generate new associations.

In the contrasting information processing account that we here consider (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*, the delay of reinforcement .

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 conditioned response onset realtive to CS onset 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)

Table 1: Different Answers to Basic Questions

• Why does the conditioned response appear?

Associative model (AM)*:* Because the associative connection gets stronger (the reinforcing effect of the US).

Computational model (CM) 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. Because there is a loss of net excitatory associative strength.

CM. 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 inter-reinforcement interval, the remembered CS-US interval, or both.

• What is the effect of non-reinforcement (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. None; 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 CS onset?

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 arise (in which the animal learns to inhibit responding to CSs that predict the withholding of the US)?

AM. The omission of an otherwise expected US (the occurrence of a No-US) strengthens an inhibitory association.

CM. The additive solution to the rate-estimation problem yields a negative rate of reinforcement for the "inhibitory" CS. This negative estimate is added to the positive estimates for the other CSs, reducing the predicted increase in the rate of reinforcement

• Why is the latency of the conditioned response proportional to the latency of reinforcement?

AM. There is no widely accepted answer to this question in associative theory.

CM. Because the animal remembers the reinforcement latency and compares the currently elapsing interval since CS onset to that remembered interval.

• What happens when a CS follows the US rather than preceding it (backward conditioning)?

AM. Nothing: only forward temporal pairing produces associations. *Or*, an inhibitory connection between CS and US is formed.

CM. A negative CS-US interval is recorded, or, equivalently, a positive US-CS interval. (remembered intervals, like remembered rates, are signed.)

• How does a secondary CS acquire behavioral potency (a secondary CS predicts a primary CS, which alone has been directly paired with the US)?

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 CS 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.

Time Scale Invariance

The value of reconceptualizing long familiar phenomena lies in the extent to which the new conceptual framework leads to experiments that deepen our insight into the fundamental nature of the phenomenon. The information processing framework brings into sharp focus a quantitative principle about conditioning, which, 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. Gallistel and Gibbon (2000) discuss 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).

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 Figure 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 Figure 2. In a partial reinforcement protocol, reinforced and unreinforced trials--trials where 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, non-reinforcement 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 9 unreinforced trials for every 1

reinforced trial (Figure 2, see also Williams, 1981).

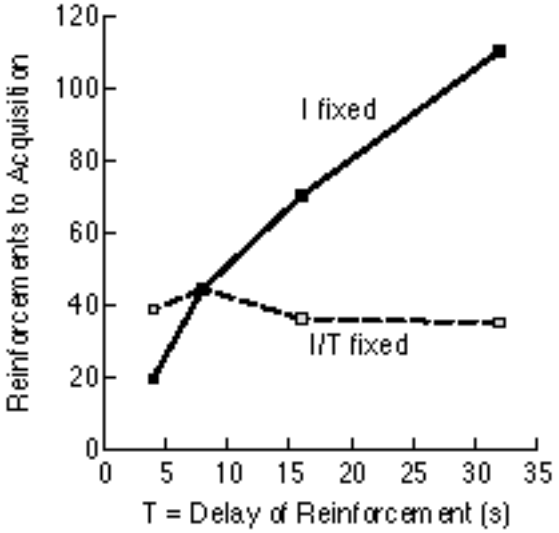

Figure 1. *Reinforcements (feedings) required for the onset of reliable conditioned responding in experiments by Gibbon et al. (1977). The experiments varied both the delay of reinforcement, T (the interval between CS onset and the US delivery, which coincides with CS offset), and the intertrial interval,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

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 the latter estimate grows in proportion to the cumulative intertrial interval. 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, 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.

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 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.) Associative theorizing also makes extensive use of intrinsically decaying stimulus traces (Wagner, 1981; Sutton & Barto, 1990). The intrinsic decay rates make the predictions of these models very sensitive to the time-scale of the protocol. Other associative models (for example, 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, 1986).

In summary, the use of a wide variety of conditioning paradigms and subject species to determine of 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 our 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, for example, Tang, et al., 1999) are not time-scale invariant processes.

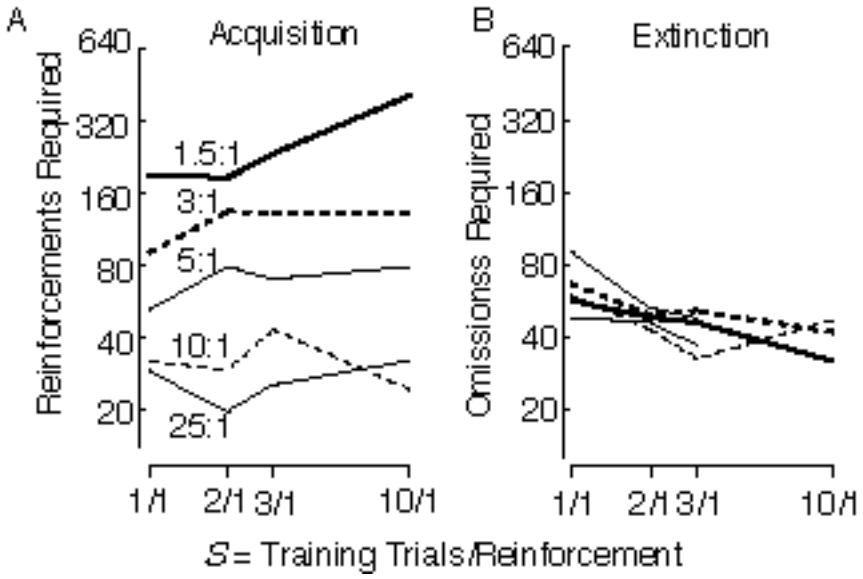

Figure 2. A. *The effects of the I/T ratio and the reinforcement schedule on reinforcements to acquisition in the Pavlovian conditioning of pigeon key pecking. (I is the intertrial interval and T the delay of reinforcement after the CS comes on). The number pairs by each curve (e.g., 5:1) give the ratio between these two intervals. The bigger this ratio the fewer the reinforcements required. By contrast, S (trials/reinforcement) has litte or no effect.* B. *The number of expected reinforcements that must be omitted to extinguish conditioned responding as a function of the I/T ratio and the schedule of reinforcement, S. Neither variable has much effect. (Replotted from data in Figure 14 of Gallistel & Gibbon, 2000.)*

References

Brown, T. H., Kairiss, E. W., & Keenan, C. L. (1990). Hebbian synapses: biophysical mechanisms and algorithms. Annual Review of Neuroscience, 13(8), 475-511.

Gallistel, C. R., & Gibbon, J. (April 2000). Time, rate and conditioning. Psychological Review.

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.

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.

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, (Vol. 20, pp. 137-192). New York: Academic.

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, (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 compuational neuroscience: Foundations of adaptive networks, (pp. 497-537). Cambridge, MA: Bradford/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: Lawrence Erlbaum.

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

Recommended Readings

Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. Psychological Review. [information processing models applied to a wide range of basic experimental results in the conditioining literature]

Gallistel, C. R. (1990). The organization of learning. Cambridge, MA: Bradford Books/MIT Press. [an information processing approach to learning in general]

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. American Psychologist, 43, 151-160. [emphasizes the many important experimental insights about conditioning that have not found their way into textbooks]