


Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Neurobiology of Learning and Memory

journal homepage: www.elsevier.com/locate/ynlme

Review

Time to rethink the neural mechanisms of learning and memory

Charles R. Gallistel^a, Peter D Balsam^{b,*}^a Center for Cognitive Science, Rutgers University, United States^b Departments of Psychology and Psychiatry, Barnard College and Columbia University, United States

ARTICLE INFO

Article history:

Available online 3 December 2013

Keywords:

Learning
Conditioning
Timing
Time perception
Anticipation
Information theory

ABSTRACT

Most studies in the neurobiology of learning assume that the underlying learning process is a pairing – dependent change in synaptic strength that requires repeated experience of events presented in close temporal contiguity. However, much learning is rapid and does not depend on temporal contiguity, which has never been precisely defined. These points are well illustrated by studies showing that the temporal relations between events are rapidly learned– even over long delays– and that this knowledge governs the form and timing of behavior. The speed with which anticipatory responses emerge in conditioning paradigms is determined by the information that cues provide about the timing of rewards. The challenge for understanding the neurobiology of learning is to understand the mechanisms in the nervous system that encode information from even a single experience, the nature of the memory mechanisms that can encode quantities such as time, and how the brain can flexibly perform computations based on this information.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

The neurobiology of learning has been guided by the idea that knowledge is acquired through associative learning. Pavlovian conditioning, the prototype of associative learning, is believed to occur because of repeated pairings of a conditioned stimulus (CS) with an unexpected unconditioned stimulus (US). A century of research has led to the accepted generalization that this learning depends on contiguity and that, in most cases, it requires many trials before it is complete. Thus, the search for the mechanisms underlying learning has focused on neural changes that depend on contiguity and repetition. There is, however, accumulating evidence that this view fails to capture a number of critical features of the learning process and fails to appreciate a fundamental function of memory. Here we highlight the shortcomings of the traditional view and sketch out an alternative information theoretic approach. We emphasize the data consistent with this approach but the reader should be aware that not all the extant data on Pavlovian conditioning are captured by this alternative. We note below when there are exceptions to the generalizations that form the foundation of this new approach.

Much of the evidence that caused us to challenge the classic view comes from studying the role of time in conditioning. Time was thought to modulate the learning of associations in the sense that temporal contiguity was necessary for learning – the less the contiguity between CS and US the weaker the resulting associative

bond and/or the more slowly it developed. The formation of the associative bond was sensitive to the temporal interval, but the bond did not encode that interval. That is, one could not recover the interval from knowledge of the strength of the association it produced, because many other factors also influenced that strength. However, it was already evident early in the study of Pavlovian conditioning that the interval between the onset of the CS and US presentation was in fact learned. As early as [Pavlov \(1927\)](#) it was known that the strength of anticipatory conditioned responses (CR's) grows during the presentation of a prolonged CS that signals a fixed delay to the US, a phenomena that Pavlov called inhibition of delay. Since those early observations of Pavlov it has come to be accepted that the learning of specific temporal intervals occurs during these protocols (see [Balsam, Drew, et al., 2010](#); [Molet & Miller, 2013](#); [Ward, Gallistel, et al., 2013](#)). As this research has progressed, it has become evident that times seem to be learned extremely rapidly, from even single experiences and even before an anticipatory CR emerges ([Drew, Zupan, et al., 2005](#); [Ohyama & Mauk, 2001](#); [Ward, Gallistel, et al., 2012](#)). A dramatic example of rapid temporal learning is presented in [Diaz-Mataix, Ruiz Martinez, et al. \(2013\)](#). In one of their experiments, rats were exposed to a Pavlovian fear conditioning procedure in which a single presentation of a tone was followed by a shock 30 s later. This was sufficient to produce reliable freezing to the tone. The next day subjects were given a reminder trial, which consisted of a few additional pairings of the tone and shock. Different groups of subjects were given the shock at the training time (30 s) or at a different time (e.g. 10 s) after the onset of the tone. In order to see if a reconsolidation process was triggered; half the subjects received

* Corresponding author.

E-mail address: balsam@columbia.edu (P.D Balsam).

an infusion of a protein synthesis inhibitor into the basal lateral amygdala following the reminder trial while the remainder of the subjects received vehicle infusions. The memory was vulnerable to disruption in only those subjects that experienced the shock at a new time. The rats had encoded the time in the original learning and a few presentations of the shock at a new time was enough to trigger an updating of the memory. Other studies show that The CS–US interval can be encoded in a single trial (Davis, Schlesinger, et al., 1989). Thus the encoding of temporal information is indeed rapid.

Given such findings, important questions for the neurobiologist to pursue are (1) what learning mechanisms are plausible given that information is encoded in a single experience; (2) how does the nervous system store information about a specific duration; and (3) How does the knowledge about time affect the expression of behavior?. We amplify the challenge that these three questions pose below. We start with the third question because these behavioral studies put important constraints on the possible answers to the first two questions.

1.1. Temporal information and the modulation of behavior

Research on the effects of varying temporal parameters in conditioning protocols casts serious doubt on the widespread belief that temporal contiguity—as ordinarily understood—is a foundational principle of learning (Balsam & Gallistel, 2009; Balsam et al., 2010; Clayton & Dickinson, 1998; Raby, Alexis, et al., 2007). First, is that learning occurs over very long delays sometimes lasting days (Clayton & Dickinson, 1998; Raby et al., 2007), as first became evident with the discovery of poison-avoidance learning (Garcia, Kimmeldorf, et al., 1961; Holder, Bermudez-Rattoni, et al., 1988). Even in standard conditioning protocols increasing CS–US intervals does not weaken learning; rather it changes how that learning is expressed. For example, if a brief presentation of a keylight is paired with grain, a pigeon will come to peck at the light – a procedure known as autoshaping (Brown & Jenkins, 1968). As would be expected from a contiguity point of view the briefer the interval from light onset to the presentation of grain, the sooner the subject comes to peck at the light. Consider, however, what happens when contiguity is changed in two different ways. First, if a light comes on and remains on for a long time before the grain, the bird does not peck at it. Instead the bird becomes hyperactive and paces back and forth in the chamber (Mustaca, Gabelli, et al., 1991; Silva & Timberlake, 2005). Thus, a long CS does not result in a failure of learning; the learning is intact, but the way it is expressed changes based on the duration of the CS (see also Holland, 1980). A second way to vary contiguity is to keep the CS duration constant but to introduce a gap between the offset of the CS and the onset of the US. This is called trace conditioning, and it is well known that when the gap gets larger CR's are weaker – in autoshaping experiments the pigeons become less likely to peck at the keylight as the trace interval is lengthened (Balsam, 1984). However, the failure to peck at the keylight is not a failure of learning. When the keylight signals the bird that it is about midway between one food and the next, the bird turns away from the light and actively retreats to a distant location (Kaplan, 1984). While there are alternative interpretations of these data (Brandon, Vogel, et al., 2003), from the perspective we present here, the bird has no trouble learning the temporal relation between the keylight and food but the behavior that is controlled by the cue is appropriate to having learned that the cue signals a long delay to the next reward (see Brandon, Vogel, et al., 2003 for an alternative interpretation.). Thus, it appears that contiguity has little impact on whether or not learning occurs, but it does have a major impact on how learning is expressed. Said another way, failures to observe anticipatory CRs should not be interpreted as failures of learning.

Another difficulty for a contiguity view of learning comes from the unsolved problem of specifying what constitutes a temporal pairing. The traditional view, rendered explicit in formal models, is that the associative process imposes a window of associability that has some intrinsic width (Gluck & Thompson, 1987; Hawkins, Kandel, et al., 2006). If the CS–US interval is less than the width of the window, an association forms between the neural elements excited by these two different stimuli. If the interval is wider than the window, no association forms. However, the width of the window has never been experimentally specified, even for a given CS (e.g., tone) and US (e.g., shock) in a given species (e.g., rat). Rescorla (1972) reviews attempts to determine the critical delay and concludes that all have failed.

The problem with the concept of a window of associability—a critical interval that defines what we understand by CS–US contiguity—goes beyond our inability to determine experimentally what that critical interval is. In the Rescorla (1968) experiments that demonstrated that contingency—not simple contiguity—governed conditioning, the US's were presented at random times. Onset of the CS did not predict a US at some fixed interval, as in delay conditioning; rather it announced a change in the rate of US occurrence. Because this rate was random, there were occasions in which the CS came and went without a US and others in which more than one US occurred during the CS. This raises the question of where in time we should imagine that the window of associability is located relative to the onset of the CS, and what happens when more than one US falls within a single window, and what happens when one falls within the window and another outside it, and so on. This problem becomes acute in the case of context conditioning. The “CS” (that is, the chamber itself) is present for many minutes and many US presentations occur at random times while it is present. In sum, despite the popular belief in contiguity, the notion that there is a critical CS–US interval has never been formulated in a way that survives empirical tests or deals with the conceptual problems raised by the variety of protocols that produce excitatory conditioning despite the lack of discrete pairings of CS and US.

1.2. Time and the emergence of anticipatory conditioned responses

The idea that temporal information drives the emergence of CRs owes its roots to the observation that the speed of conditioning in autoshaping depends on the ratio of the time between US's, referred to as the cycle time (C), and the duration of the CS–US interval, referred to as the trial time (T). Across a broad range of values, the number of trials to acquisition is determined by the C/T ratios, regardless of the absolute values of C and T (Gibbon, Baldock, et al., 1977; Ward et al., 2012). This relation is illustrated in Fig. 1a which shows the results of autoshaping experiments (Gibbon et al., 1977) where different groups of subjects were exposed to protocols that differed with respect to the duration of the CS–US interval. In groups for which the ITI was held constant, the trials to acquisition increased with increasing CS–US interval. However, in groups for which the ITI was increased proportionally to increase in the CS–US interval, the number of trials to acquisition was constant. Thus, what matters in terms of the associability of a CS with a US (speed of conditioning) is not the CS–US interval per se but rather the proportion this interval bears to the US–US interval. While the speed of CR emergence is determined by the C/T ratio in these autoshaping experiments, it is not yet clear that this is true for all conditioning preparations. Across a moderate range of values it is true for appetitive head-poking in rodents (Ward et al., 2012) but may break down with very long CS durations (Holland, 2000; Lattal, 1999). In aversive conditioning the degree of suppression produced by a CS associated with shock is determined by the C/T ratio (Coleman, Hemmes, et al., 1986; Stein, Sidman, et al., 1958) but




Fig. 1. Acquisition speed as a function of trial CS duration. Different groups of pigeons were exposed to autoshaping protocols in which trial CS durations (T) ranged from 4 s to 32 s across groups. For some groups, the duration of the intertrial interval was kept constant (ITI fixed; filled circles). Panel A shows that for these groups, the number of trials to acquisition increased with increased CS duration. In the other groups the inter trial interval was increased proportionally to the increase in CS duration so the ratio of the two intervals was kept constant (ITI/ T fixed; filled squares). In these groups, speed of acquisition remained constant regardless of the duration of the trial CS. After Balsam et al. (2010). Original data from Gibbon et al. (1977). Panel B shows how the temporal informativeness of the CS changed in the different protocols. The speed of acquisition shown in panel A is directly related to the mutual temporal information between the CS and US which is shown in panel B (after Ward et al., 2013).

to our knowledge the effects of C/T ratios on fear conditioning and eyelid conditioning have not been directly studied. As noted above, the form of conditioned responses changes with alterations in temporal parameters. In autoshaping, when CS's become long birds will increase general activity rather than peck at a CS. Similar changes in the topography of CR's have been described in other species (Akins, Domjan, et al., 1994; Silva & Timberlake, 1997, 1998, 2005). Thus the impact of changing the C/T ratio is on the likelihood of a particular CR topography. Consequently, changes in the likelihood of a particular CR should not necessarily be taken as a reflection of whether or not learning has occurred.

We digress in this paragraph to call neurobiologists' attention to the methodological importance of distinguishing between the parameter-trade-off method used by Gibbon, Baldock, et al. (1977) and Ward et al. (2012) and the more common "psychometric" method for studying the effects of protocol variables in learning. The parameter-trade-off method determines combinations of experiential parameters that yield the same behavioral result. The Gibbon et al. experiment varied both the CS-US interval, the US-US interval and measured the number of training trials required to cause the same behavioral result, namely, the appearance of a conditioned response. The parameter-trade-off method is standard in sensory psychophysics, where, for example, one determines a spectral sensitivity function by finding the combinations of wavelength and light intensity that produce the same frequencies of seeing. We believe that this approach provides better leverage for those interested in the neurobiology of learning than the more traditional approach of looking to a learning curve for quantitative information about the underlying process, as is widely done in studies of the neurobiology of learning and memory. The learning-curve approach is equivalent to the determination of a psychometric function in psychophysics. It plots a measure of performance against a "stimulus" variable, namely, the number of "trials" (CS-US pairings). The equivalent function in visual research is the plot of the frequency of seeing against, say, light intensity. Psychometric functions rarely reveal important quantitative properties of underlying physiological mechanisms because they confound many performance-relevant variables. The parameter-trade-off method has, by contrast, yielded a rich body of quantitative information about molecular and cellular level processes in vision and elsewhere in the history of behavioral neurobiology. For example, the behaviorally determined scotopic spectral sensitivity curve, which was first determined in the 19th century, reveals the absorption spectrum of rhodopsin. It gave us this critical information before we knew

there was such a thing as rhodopsin. For additional explanation of the power of the parameter-trade-off methodology see Gallistel, Shizgal, et al. (1981).

As we noted above, the average inter-event intervals are learned rapidly, even before the CR emerges. In many preparations, knowledge of the expected time to the US guides responding even as CR's first emerge (Balsam, Drew, et al., 2002; Drew et al., 2005; Kirkpatrick & Church, 2000; Ohyama & Mauk, 2001; Ward et al., 2012). Even when the temporal control of the CR is slower to emerge (e.g. Delamater & Holland, 2008; Rescorla, 1967), one should not take its slow emergence as strong evidence against the rapid learning of time. The prior history of feeder training may establish a steady level of checking the feeder in the early stages of appetitive conditioning and some responses such as freezing may initially have low thresholds that change slowly to reflect the expected time of an anticipated aversive stimulus. Some experiments have shown that even when CR timing is initially poor, transfer tests show that times were indeed learned in just one or two trials (Davis et al., 1989; Diaz-Mataix et al., 2013). Furthermore, blocking and overshadowing are greatest when the cues of the compound share a common CS-US interval (Molet & Miller, 2013). In other words, the durations of the intervals between events are a key aspect of the content of learning in conditioning protocols, along with, of course, the encoding of the sensory properties that enable the brain to distinguish one CS from another and one US from another. A neurobiological understanding of the conditioning process will require identifying the mechanism by which quantities, such as interval durations, are encoded into enduring changes at some level of molecular or cellular neuronal structure. And, it will require an understanding of the computational mechanisms that translate this stored information into anticipatory behavior. It is generally assumed that the structural change that encodes information is at the level of the synapse and that the computational mechanisms are implemented at the level of neural circuitry. These long-standing presuppositions should not be taken as established facts. Intracellular mechanisms realized at the molecular level are also conceivable. Micro RNAs, for example, have a digital structure that suits them for information storage. They could also implement computational operations on stored information, as clearly occurs in the mechanisms that translate the hereditary information stored in DNA into organic structure and organic process. So long as we remain ignorant of the mechanisms of information storage and computation in nervous tissue, we should keep an open mind as to the structural level at which these processes occur.

2. Temporal informativeness in learning the relation between stimuli, responses and rewards

Finding the neurobiological basis of learning and memory is a challenging problem. Consideration of successful examples of material reduction from the history of behavioral neuroscience and genetics may help to direct our efforts. Examples of such successes are proving that the action potential is the nerve impulse and proving that hereditary information is stored in the base-pair sequences along the double-helical DNA molecule in a chromosome. Both of these successful efforts (and many others) have depended on measuring by indirect methods quantities that could be measured by direct methods if and when one had an hypothesis about the physical identity of the mechanism whose properties one measured indirectly. For example, Helmholtz (Helmholtz, 1852) measured the conduction velocity of the nerve impulse by the difference-in-reaction time method, which continues to be a staple of psychology, particularly cognitive psychology. At about the same time, his friend, du Bois-Reymond, discovered the (compound) action potential and hypothesized that it was the physical realization of the nerve impulse (du Bois-Reymond, 1848). A key test of his hypothesis was passed when it was shown that the conduction velocity of the action potential was the same as the conduction velocity of the nerve impulse (see Hermann, 1879 for review). For further example, classical geneticists routinely measure the linear order of non-independently assorting genes (genes on the same chromosome) by studying patterns of trait assortment. They still do this, but they also did it in the days before one knew that the double helix was the molecular structure whose properties were being measured. They did these measurements even before we knew that the hereditary material, whatever it was, was in the chromosomes. These classical genetic measurements played a fundamental role in establishing the fact that the hereditary material is in the chromosomes and that it is the DNA in the chromosomes that is the physical realization of the gene (Judson, 1980). For a final example we reiterate (see Section 1.2) that the behavioral determination of the scotopic spectral sensitivity curve played a central role in establishing that rhodopsin is the key transduction molecule in rods. Because the absorption spectrum is a molecular signature, demonstrating that the absorption spectrum of rhodopsin matched the scotopic spectral sensitivity function was very nearly all the proof anyone could want. There are many other proofs but they all depend on the correspondence between indirect psychophysical measurements and direct physical measurements made on anatomical and biochemical structures.

Our ability to identify the biophysical basis of learning and memory would be enhanced if we could find something that we can measure now, without knowing what the neurobiological mechanism of memory is, and that we may also measure by physical methods, when we have promising hypotheses about the physical changes by which brain tissue encodes experiential facts, such as interval durations. When we believed that there was a critical interval in associative learning, measuring that interval behaviorally was a promising route to identifying the neurobiological mechanism of association formation (Abrams & Kandel, 1988; Carew, Walters, et al., 1981; Gluck & Thompson, 1987). The hopes invested in that route would seem to be misplaced if we are correct that there really is no critical interval for learning. As described above, our rejection of the simple associative view rests on the observation that the conditioning depends on the C/T ratio not on the duration of T (Gibbon & Balsam, 1981). Those same experimental findings, however, have led us to a new and more promising measure of the sought-for kind, the mutual information between CS timing and US timing.

Mutual information is a statistical quantity that can be measured even when the entities or processes between which the information is shared are of fundamentally different natures. For example, one can measure the mutual information between a spike train and the overall pattern of movement of the visual field of the housefly (Rieke, Warland, et al., 1997). Thus, there is no reason why one cannot compare the amount of information encoded in a structural change in some aspect of nervous structure—whether in a pattern of changes in synaptic conductances or in changes in micro RNA sequences—and compare that direct physical measure to the measure of the mutual information in a conditioning protocol. If we could show that the information in the altered physical structure was the same as the mutual information between CS and US that produced that alteration in structure, that would be strong evidence that we had finally found the information-carrying structure – the neurobiological basis of memory. This evidence would be analogous to the evidence produced by measuring the velocity of propagation of an action potential and showing that it matches the velocity of propagation of the nerve impulse in the same nerve.

Measuring the mutual information in conditioning protocols suggests itself because the Gibbon and Balsam (1981) findings on the critical role of the C/T ratio in Pavlovian conditioning suggest that mutual information is the experience-derived quantity in the brain that determines the CS–US “associability” in a given protocol. The CS–US associability is the rapidity with which a conditioned response appears (the inverse of trials to acquisition). The Gibbon and Balsam (1981) empirical generalization implies that CS–US associability is determined by an easily measured objective component of the mutual temporal information between the onset of the CS and the onset of the US (Balsam & Gallistel, 2009; Ward et al., 2012). That in turn suggests that the experience provided by a protocol instills an encoding of that quantity somewhere in the animal’s brain and that it is this structural change that leads to the appearance of a conditioned response.

This interpretation of the Gibbon and Balsam (1981) result in terms of (a component of) the mutual information between CS and US leads to an unexpected unification in our understanding of the role of temporal parameters, stochastic parameters (partial reinforcement), cue competition and contingency in Pavlovian and instrumental conditioning (Balsam et al., 2010). These are all topics of fundamental importance to our understanding of conditioning, but they have been experimentally analyzed and theorized about as quite separate problems, often treated in separate chapters in textbooks. A partial exception to this generalization is that it has been widely understood since Rescorla (1968) that “contingency,” if it could be defined and measured, was somehow closely connected to cue competition in Pavlovian conditioning and the assignment of credit problem in instrumental conditioning. Information theory provides a generally applicable measure of contingency (Gallistel, 2012a, 2012b; Gallistel, Craig, et al., 2013).

The temporal information that the onset of a CS provides about the onset of the US is the reduction in the uncertainty about when to expect the next US. This reduction is the mutual temporal information between CS and US. How this mutual temporal information is most easily measured varies from protocol to protocol. In the “truly-random-control” experiment (Rescorla, 1968), the US is generated by a random rate (Poisson) process whose parameter may or may not increase at CS onset and decrease at CS offset. The measure of mutual temporal information is the difference in the entropies of the distribution of US–US intervals in the presence and absence of the CS. If we measure information in bits, the mutual temporal information between CS and US in this very simple protocol is $\log_2(\lambda_{CS}/\lambda_{-CS})$, where λ is the rate parameter of the Poisson process (see Balsam and Gallistel (2009) for derivation).

From the formula just given for computing the mutual information in the protocol where the rate of US occurrence varies depending on whether the CS is or is not present, we see immediately that when the two rates are equal, as they were in Rescorla's (1968) truly random control condition, then their ratio is 1, in which case the mutual information [$\log(1)$] is 0. This explains why the rats in that condition did not respond to the CS during final testing despite the fact that the CS had been paired with the US just as frequently in that control condition as in the other conditions, where rats did respond to the CS. It was this result that led to the suggestion that it is contingency, not temporal pairing, that drives conditioning (Rescorla, 1972). This result is a foundational result in the literature on cue competition, so we begin to see how a focus on the mutual information between CS and US unifies our understanding of Gibbon and Balsam (1981) with our understanding the other basic findings in the voluminous cue-competition literature—blocking and overshadowing (Kamin, 1969a, 1969b) and relative validity (Wagner, Logan, et al., 1968). All of these findings are predicted if one computes the mutual information between a given CS and the US after subtracting out the information about the timing of the next US that comes from the other CSs (and from the US itself, when the US–US interval is fixed). In other words, one computes the reduction in the residual uncertainty that occurs at CS onset. The residual uncertainty is the uncertainty that remains after other sources of information about the timing of the US have been taken into account. We suggest that the brain makes this same computation, which is why we see such clever behavioral adaptations to complex predictive relationships reported in the cue competition literature.

In the more complex delay conditioning protocol, which is widely used in neurobiologically oriented work, the mutual information may be decomposed into three components, two of them “objective” and one “subjective”. One objective component depends only on the C/T ratio (see Balsam and Gallistel (2009) for derivation). We have already seen that the rapidity with which a conditioned response appears depends on this parametric property of the protocol. Fig. 1B formally shows that the speed of acquisition is monotonically related to this component of the mutual information. When the C/T ratio is constant so is acquisition speed, but as this information is increased, acquisition is faster.

The other objective component of the mutual information between CS and US depends on the partial reinforcement schedule (see Gallistel, 2012b for derivation). Computing the contribution of this component predicts two results from the literature on the effects of partial reinforcement that are profoundly counterintuitive and hard to explain when conditioning is viewed from an associative perspective (Gallistel, 2012a, 2012b). The first prediction is that the number of reinforced CS presentations required for the appearance of the conditioned response is not increased by interpolating large numbers of unreinforced CS presentations (Gibbon, Farrell, et al., 1980; Gottlieb, 2004, 2005; Harris, 2011; Williams, 1981). For exceptions to this generalization, see Bouton and Sunsay (2003), Gottlieb and Rescorla (2010). This is profoundly puzzling from an associative perspective, because the interpolated unreinforced presentations of the CS should weaken the CS–US association and thereby retard the growth of associative strength. On the other hand, a simple information-theoretic computation shows that halving the proportion of reinforced trials doubles the information communicated by the remaining reinforced trials, thereby explaining why in many cases partial reinforcement does not increase reinforced trials to acquisition (Gallistel, 2012a, 2012b).

An information-theoretic analysis also predicts that partial reinforcement should increase the number of unreinforced trials required to eliminate the conditioned response during extinction, when CS presentations are no longer reinforced. Indeed, partial reinforcement during the initial conditioning phase increases

trials-to-extinction, and it does so in proportion to the thinning of the reinforcement schedule (Gibbon et al., 1980). From an associative perspective, partial reinforcement during conditioning should decrease trials to extinction, not increase it. This puzzle is called the partial reinforcement extinction effect. It was observed long ago that, “The most critical problem facing any theory of extinction is to explain the effect of partial reinforcement.”—(Kimble, 1961, p. 286). Theorists have proposed several ways that associative theory might accommodate something like this result, although, to our knowledge, none that predicts that the increase in trials to extinction scales with the thinning of the reinforcement schedule. But, simple information-theoretic calculations show that the partial reinforcement during training decreases the per-trial rate at which information that there has been a decrease in the schedule of reinforcement accumulates during the extinction phase. Whether extinction is based on a change in the rate of reward (Gallistel, 2012a, 2012b; Gallistel & Gibbon, 2000) or a change in the per trial likelihood of reward (Drew, Yang, et al., 2004; Gallistel, 2012a, 2012b; Haselgrove, Aydin, et al., 2004), the decrease in the per-trial rate of information accumulation during the extinction phase is proportional to the thinning of the reinforcement schedule during the conditioning phase. Thus, halving the schedule of reinforcement during the conditioning phase doubles the number of extinction trials required to give the same amount of information about the change (Gallistel, 2012b, but see Gottlieb & Prince, 2012 for discussion of conditions where this generalization may not hold). This explains the cases in which the effect of partial reinforcement on extinction is to increase trials to extinction in proportion to the thinning of the reinforcement schedule during the conditioning phase. Further work will be required to understand factors that might contribute to the failure to find scaling in all cases. In particular, understanding exactly how uncertainty about *when* something will occur combines with uncertainty about *whether* it will occur at all will be central to generalizing the approach we present here.

The third component of the mutual information between CS and US in a delay protocol is called the subjective component because it depends only on the precision with which the subject represents intervals (see Balsam & Gallistel, 2009; Balsam et al., 2010 for derivations). In other words, the amount of information in this third component depends only on the subject's Weber fraction for time, a measure of the relative precision with which it represents durations. The contributions of the other two components depend only on parameters of the protocol (the C/T ratio and the partial reinforcement schedule), which is why we call them the objective components of the mutual information. Ward et al. show that this third component does not affect trials to acquisition. This finding makes sense in that it means that the co-variation between the number of trials to acquisition and the parameters of the protocol depends only on those parameters (the structure of events in the world), not on a property of the animal.

Finally, the measurement of mutual information gives us a measure of contingency, namely the ratio between the mutual information and the basal US entropy (Gallistel, 2012a, 2012b). The basal entropy is the baseline uncertainty about when the next US will occur. This is the entropy of the distribution of US–US intervals after convolution with the precision with which the subject's brain represents the durations of intervals. The convolution with the brain's precision of interval representation is necessary because, when the US–US interval is fixed, the objective distribution is the Dirac delta function, which has 0 entropy. Intuitively, when the US–US interval is fixed, your uncertainty about when the next US will occur, given that you know when the last one occurred, is limited only by the precision with which you can represent the fixed US–US interval. If you could represent it perfectly, you would have no uncertainty about when to expect the next US, or about when to expect any future US, no matter how remote.

The information-theoretic measure of contingency also suggests a solution to the assignment-of-credit problem in instrumental conditioning (Staddon & Zhang, 1991; Sutton, 1984). This is the problem of deciding which previous actions are responsible for generating reinforcements. Put another way, how does the brain determine which responses produce which outcomes? An information-theoretic solution to this problem is to assign credit to previous actions (and events) on the basis of the retrospective temporal contingency between the reinforcements and the past actions and events. A retrospective contingency exists between a response, R , and an outcome, O , when the entropy of the distribution of $R \leftarrow O$ intervals (intervals measured from O back to the most recent instance of R) is less than the entropy of the distribution of $R \leftarrow R$ intervals. A prospective contingency exists when the distribution of $R \rightarrow O$ intervals has less entropy than the distribution of $O \leftarrow O$ intervals.

Prospective and retrospective contingency can be grossly asymmetrical. When, for example, a subject responds on a partial reinforcement schedule (of whatever nature), the prospective contingency is weak, because the distribution of $R \rightarrow O$ intervals typically has an entropy only moderately less than the entropy of the $O \leftarrow O$ distribution. This obtains because many R s do not trigger an O , so the intervals from those R s to the next O are almost as long and variable as the $O \leftarrow O$ intervals themselves. By contrast, the distribution of $R \leftarrow O$ intervals has essentially 0 entropy, because every O is preceded by an R at a very short fixed interval (on the order of 0.1 s). In other words, given the time at which an R occurred, there is great uncertainty about when the next O occurred, but given the time at which an O occurred, there is negligible uncertainty about when the immediately preceding R occurred. Put yet another way, the time at which an outcome occurs is retrodictively almost perfectly informative about the time at which an immediately preceding response occurred, but the time at which an R occurs is only weakly informative about when the next O will occur.

Retrodictive competition (credit assignment) works much like predictive competition (cue competition): If for a given O , there is a single response event, R_1 , such that the distribution of $R_1 \leftarrow O$ intervals has no entropy, then there is, so to speak, nothing left to explain as regards what causes O ; R_1 gets all the credit. If, on the other hand, either R_1 (for example, pressing the lever) or R_2 (pulling the chain) causes O , then the distribution of $R_1 \leftarrow O$ intervals and the distribution of $R_2 \leftarrow O$ intervals will both have substantial entropy, but the distribution of $[R_1|R_2] \leftarrow O$ intervals will have negligible entropy. In that case, a more complex model explains everything that there is to explain in regard to the causation of this O . On this hypothesis, the brain solves the assignment of credit problem by entertaining increasingly complex causal models until there is nothing left to explain or until increases in model complexity do not reduce the residual retrospective uncertainty (do not increase explanatory power).

The information-theoretic account of cue competition and credit assignment explains the success of the Rescorla–Wagner model of conditioning. The RW model assumes that associative strengths add, and it assumes that there is an upper limit on the sum. These are postulated properties of the associative process. They do not reflect objective (non-psychological) facts about the experienced world. On the information-theoretic account, these assumptions about the associative process reflect objective, mathematical truths deeply relevant to the construction of a useful model of event dependencies: independent entropies add and the source entropy (otherwise known as the available information) is the upper limit on the amount of information that all predictors (or retrodictors) combined can convey.

The relevance of the information-theoretic analysis of conditioning phenomena to the neurobiology of learning depends on

the stance one adopts toward its successes. One stance is to regard it as a mathematically powerful unifying summary of a wide range of phenomena couched in abstractions that have nothing to do with the underlying neurobiological mechanisms. On this view, the animal behaves “as if” its brain computed the entropies of distributions and mutual information, but when we someday understand the underlying neurobiological mechanism we will see that this is not really what it is doing. For many years, this was the stance that many biochemists adopted toward genetic theory (Judson, 1980): The mechanism by which species-specific form and variants therein (traits) were transmitted from generation to generation behaved as if there were several different sets of linearly arranged information-carrying “particles” (genes) that somehow made copies of themselves and somehow encoded species-typical form and variants thereof, but, because it was impossible to imagine what such a mysterious thing might look like at a biochemical level of analysis, many biochemists believed genes were a convenient fiction, to be supplanted someday by a biochemically intelligible mechanism. Fortunately, this was not the stance of Watson and Crick and other important figures in the history of molecular biology. They believed in the material reality of the gene. We believe in the material reality of the information-theoretic analysis of conditioning. We believe that the neural mechanism of learning explicitly (symbolically) represents distributions and their entropies and does the computations on these entropies that define mutual information and contingency.

3. Challenges for the neurobiology of learning

Given the success that we have had in understanding behavior as guided by the extraction of temporal information from the continuous stream of experience, we think that the neurobiologist is challenged to answer several related questions. (1) What learning mechanisms are plausible given that information is encoded in a single experience?; (2) How does the nervous system store information about a specific duration?; (3) How does the nervous system compute information from the stored experience?

3.1. What learning mechanisms are plausible given that time is encoded in a single experience?

Most models of timing assume that there is an interval timer used to measure the experience of intervals ranging from a few seconds to hours or days (for reviews Buhusi & Meck, 2005; Gallistel & King, 2009; Lustig, Matell, et al., 2005; Matell & Meck, 2004). There is an important conceptual and practical (implementational) problem with this assumption. The first time an event occurs marking the onset of an interval that turns out to be behaviorally important, there is no way of knowing whether to start an interval timer or not, nor how many to start, because there is no foreknowledge of what may or may not follow. The possibilities are unlimited; therefore the system should start an infinite number of new timers running every time anything happens. The same start event may pair with arbitrarily many different end events, each pair delimiting a different duration. For example, the onset of a CS (CS1on, a start event) may be followed by its offset, thus delimiting a CS1on–CS1off interval, and also by the onset of a sweet milk US after a different interval, thus delimiting a CS1on–US1on interval, and also by the onset of a different CS, thus delimiting a CS1on–CS2on interval, and so on ad infinitum. If the CS1on–CS1off interval is measured by a timer that stops and resets at CS1off, then that timer cannot be used to measure any of the other intervals.

Speaking more technically, the number of timers required is exponential in the number of events that might mark the starts and ends of intervals, which any computer scientist will recognize

as a fatal implementational problem, no matter what the architecture of the computer. However if the system does not time intervals on first encounter, then there is no record of them, in which case every subsequent encounter is effectively a first encounter. Yet as we have noted in Section 1.2, the extraction of duration from a single experience has been experimentally demonstrated (Balsam et al., 2010; Diaz-Mataix et al., 2013). Thus, if the neural machinery could not time and remember a single instance of an interval, then it is unclear how it could come to know that the interval recurs (Diaz-Mataix et al., 2013). To note a recurrence, a machine, whether made of neurons or not, must have a record of a previous occurrence. Such a record is a precondition for the machine's recognizing a recurrence. More generally, recognition presupposes that one can store individual instances of whatever it is that is recognized. Thus, to recognize that the n th occurrence is the same as a previous occurrence or all previous occurrences, the machine must have made recognizable records of the previous occurrences, one by one, as they occurred.

An alternative to measuring interval durations is to compute them. If the time of each event is automatically encoded, then the interval between any two events may be computed retrospectively by subtracting the earlier time from the later time (Gallistel, 1990). This is the sort of computation involved in your judgment of how long it has been since you started reading this article. Many species are capable of retrospective temporal computation (Molet & Miller, 2013). We cannot point to a specific mechanism but times of occurrence may be obtained from the phases of the many different neural oscillators known to be present not only in neural tissue but throughout the body (Silver, Balsam, et al., 2011). Ultradian, circadian and infradian cycles are present even in bacteria and plants, so plausible signal sources are at hand. What is hard to envisage in our current state of ignorance is the mechanism(s) by which the phases of these oscillations may be read off and recorded in memory.

3.2. How does the nervous system store information about a specific duration?

When a duration has been either measured or computed, it is stored in memory and available to guide future behavior. But, we do not have an understanding of how the nervous system stores information gleaned from experience in a computationally accessible form (Gallistel & King, 2009; Gallistel & Matzel, 2013). Neurobiologically motivated theories of timing only try to explain the fact that conditioned behavior is timed. They do this by assuming that individual neurons or constellations of neurons that become active at the appropriate time become associated with the anticipatory response. In other words, they are S–R theories. For example, some associative theories posit the existence of interoceptive microstimuli triggered by CS presentation that decay with the passage of time (e.g. Brandon et al., 2003; Sutton, 1988). Some of the microstimuli are said to be contiguous with the US and these dynamically changing microstimuli can then account for why the CR might become more likely with the passage of time. Like all such theories, they are anti-representational: the brain responds as if it knew the interval, but it does not have that information stored in symbolic form, which is to say, in a form that makes it useable in computational operations. Insofar as information about the durations of the experienced intervals can be said to reside somewhere in these conditioned reflex arcs, it resides in the intrinsic dynamics of the neurons (Buonomano & Laje, 2010; Laje, Cheng, et al., 2011). There is no explanation of how any other part of the brain could have access to those dynamics. Because microstimulus models do not posit the symbolic representation and storage of the information gained from experience, they do not attempt to explain the fact that acquired

temporal information—like other kinds of acquired information— informs behavior in different and flexible ways (Arceidiano & Miller, 2002; Denniston, Blaisdell, et al., 1998; Denniston, Blaisdell, et al., 2004).

Perhaps the biggest obstacle to neurobiologists' acceptance of the view that the brain stores information in symbolic form, just as does a computer, is our inability to imagine what this story might look like at the cellular and molecular level. We hope the following observations may go some way toward overcoming this obstacle.

The storage of information in readable symbolic form is not foreign to the conceptual framework of contemporary biologists. Both DNA and RNA carry information forward in time, and there is elaborate machinery at the cellular level for reading that information. More generally, any biophysical mechanism with the properties of a readable switch can serve as a mechanism of information storage. A readable switch is something that has more than one (but usually a small number) of thermodynamically stable configurations and whose differences in configuration are causally effective, that is, they alter the course of some process. The switches in DNA and RNA are the base-pair loci; there are 4 possible configurations (possible positions of the switch) at each locus; thus, each locus can store 2 bits of information. There are other switch molecules. The opsin molecules in the outer segments of photoreceptors are photolabile switches. They have two thermodynamically stable configurations (isomers). The absorption of a photon throws the switch; it isomerizes a rhodopsin molecule. The switch configuration of the switch is readable, because it is enzymatically active only in its isomerized configuration. Methylation is another example of a switch mechanism; typically, it converts a target (e.g., a gene) from the active to the inactive state. It is one of the principal mechanisms for the determination of cell fate, a thermodynamically permanent switching process crucial to the process of tissue development.

Thus, what we must look for neurobiologically is a system of switches. Some questions that must be answered beyond the question of the biophysical identity of the switches: (1) How do neural signals impinging on receptor molecules in postsynaptic membranes throw the switches (write to memory)? (2) What is the switch code, that is, how does the configuration of the switches encode, for example, the duration of an interval? (3) How does the write process map from the dynamic code, the code by which information is carried in evanescent spike trains, to the static code, the code by which the same information is carried by the static configuration of the thermodynamically stable switches? The answer to this question will depend in no small measure on the answer to the question how spike trains convey information? Is it simply by the rate of firing (spikes per unit time), and, if so, what is the read interval (the interval over which the rate is measured)? Or is it the interval between spikes? Or is it some as still unimagined scheme? This is an unsettled fundamental question in neuroscience (cf Rieke et al., 1997) (4) How is memory read, that is, how does the configuration of switches alter the input–output properties of the neuron or neural circuit (that is, of the memory element) so as to reconstitute a neural signal that carries the information that has been written into the switches?

As our illustrative examples given above tend to imply, our guess is that the answers to these questions lie at the level of molecular changes within neurons rather than at the level of synaptic changes within neural circuits. We believe neuroscientists should consider the possibility that the neurobiological elements of memory are intraneuronal molecular structures, not neural circuits (constituted, at a minimum, of a presynaptic and a postsynaptic neuron). Among other things, this possibility increases the plausible density of information storage in the brain by many orders of magnitude.

In our experience, a first objection to the proposal that memory resides in intraneuronal molecular structures is often that intraneuronal molecular structure cannot inform behavior on a sufficiently short time scale. This intuition is demonstrably false: Pattern vision depends on the photo-isomerization of the opsins embedded in an internal membrane in the outer segments of rods and cones. The conversion of these intracellular changes in the configuration of individual molecules into neural signals and the conversion of those signals into coordinated muscular action is so fast that (some) people can hit a tennis ball struck from 25 m away and travelling at close to 50 m/s. Although we cannot point to a specific neurobiological mechanism that encodes quantities, time dependent changes in intracellular signaling pathways and alterations in thresholds based on experience seem plausible candidates for the encoding of information about specific durations.

3.3. How does the nervous system compute information from the stored experience?

Perhaps, the least understood aspect of the neurobiology of timing is how the brain can take the information stored from specific experiences and use it in flexible ways. How, for example, can the brain retrospectively compute an interval from two occurrence times, or, how, for further example, can it add and subtract remembered or computed intervals? In planning a day we take into account the sums and differences in our remembered durations of multiple activities. In sequencing actions we take into account the temporal structure of goal availability. While current studies of neurobiology may point to areas of the brain that are part of networks doing these computations the physical mechanism by which these calculations are implemented is completely unknown. Despite much diverse theorizing, we do not know how the brain stores information nor how it computes. When we do know that, we will be in a much better position to entertain neurobiologically plausible accounts of how it learns.

4. Conclusions

We are left with some fascinating unsolved problems in the neurobiology of learning. We must understand in much greater detail several steps in the learning and memory of temporal information. First, we must understand the plastic changes that underlie the rapid and encoding and storage of time (or any aspect of an event). This process must be occurring continuously. Next we must understand how patterns of events are detected and the encoding that underlies the retrieval of this information in an episodic chunk. And, finally we must understand the mechanisms of computation that allow the extraction of information that guides the decision underlying what to do at any particular time. These are all difficult challenges but ones that must be met if we are to understand the neurobiology of learning and memory.

Acknowledgments

This work was supported by was supported by National Institute of Mental Health Grant R01MH068073 to PDB and R01MH077027 to CRG.

References

Abrams, T. W., & Kandel, E. R. (1988). Is contiguity detection in classical conditioning a system or a cellular property? Learning in *Aplysia* suggests a possible molecular site. *TINS*, 11(4), 128–135.

Akins, C. K., Domjan, M., et al. (1994). Topography of sexually conditioned behavior in male Japanese quail (*Coturnix japonica*) depends on the CS-US interval. *Journal of Experimental Psychology: Animal Behavior Processes*, 20(2), 199–209.

Arcediano, F., & Miller, R. R. (2002). Some constraints for models of timing: a temporal coding hypothesis perspective. *Learning and Motivation*, 33, 105–123.

Balsam, P. (1984). Relative time in trace conditioning. *Annals of the New York Academy of Sciences*, 423, 211–227.

Balsam, P. D., Drew, M. R., et al. (2002). Timing at the start of associative learning. *Learning and Motivation*, 33(1), 141–155.

Balsam, P. D., Drew, M. R., et al. (2010). Time and associative learning. *Comparative Cognition and Behavior Review*, 5, 1–22.

Balsam, P. D., & Gallistel, C. R. (2009). Temporal maps and informativeness in associative learning. *Trends in Neurosciences*, 32(2), 73–78.

Bouton, M. E., & Sunsay, C. (2003). Importance of trials versus accumulating time across trials in partially reinforced appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 29(1), 62–77.

Brandon, S., Vogel, E., et al. (2003). Stimulus representation in SOP: I. Theoretical rationalization and some implications. *Behavioural Processes*, 62(1), 5–25.

Brown, P. L., & Jenkins, H. M. (1968). Auto-shaping of the pigeon's key-peck. *Journal of the Experimental Analysis of Behavior*, 11(1), 1–8.

Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience*, 6(10), 755–765.

Buonomano, D. V., & Laje, R. (2010). Population clocks: motor timing with neural dynamics. *Trends in Cognitive Sciences*, 14(12), 520–527.

Carew, T. J., Walters, E. T., et al. (1981). Associative learning in *Aplysia*: cellular correlates supporting a conditioned fear hypothesis. *Science*, 211, 501–504.

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, 395(6699), 272–274.

Coleman, D. A., Hemmes, N. S., et al. (1986). Relative durations of conditioned-stimulus and intertrial interval in conditioned suppression. *Journal of the Experimental Analysis of Behavior*, 46(1), 51–66.

Davis, M., Schlesinger, L. S., et al. (1989). Temporal specificity of fear conditioning – effects of different conditioned-stimulus – unconditioned stimulus intervals on the fear-potentiated startle effect. *Journal of Experimental Psychology – Animal Behavior Processes*, 15(4), 295–310.

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(2), 202–222.

Denniston, J. C., Blaisdell, A. P., et al. (1998). Temporal coding affects transfer of serial and simultaneous inhibitors. *Animal Learning and Behavior*, 26(3), 336–350.

Denniston, J. C., Blaisdell, A. P., et al. (2004). Temporal coding in conditioned inhibition: analysis of associative structure of inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, 30, 190–202.

Diaz-Mataix, L., Ruiz Martinez, R. C., et al. (2013). Detection of a temporal error triggers reconsolidation of amygdala-dependent memories. *Current Biology*, 23(6), 467–472.

Drew, M. R., Yang, C., et al. (2004). Temporal specificity of extinction in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, 30(3), 163–176.

Drew, M. R., Zupan, B., et al. (2005). Temporal control of conditioned responding in goldfish. *Journal of Experimental Psychology: Animal Behavior Processes*, 31(1), 31–39.

du Bois-Reymond, E. (1848). *Untersuchungen über thierische Elektrizität*. Berlin: Reimer.

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

Gallistel, C. R. (2012a). Contingency in learning. In N. M. Seel (Ed.), *Encyclopedia of the sciences of learning*. New York: Springer.

Gallistel, C. R. (2012b). Extinction from a rationalist perspective. *Behavioural Processes*, 90, 66–88.

Gallistel, C. R., & King, A. P. (2009). *Memory and the computational brain: Why cognitive science will transform neuroscience*. Wiley.

Gallistel, C. R., Craig, A. R., et al. (2013). Temporal contingency. *Behavioural Processes*.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, 107(2), 289–344.

Gallistel, C. R., & Matzel, L. D. (2013). The neuroscience of learning: beyond the Hebbian synapse. *Annual Review of Psychology*, 64, 169–200.

Gallistel, C. R., Shizgal, P., et al. (1981). A portrait of the substrate for self-stimulation. *Psychological Review*, 88(3), 228–273.

Garcia, J., Kimmeldorf, D. J., et al. (1961). The use of ionizing radiation as a motivating stimulus. *Psychological Review*, 68, 383–385.

Gibbon, J., Baldock, M. D., et al. (1977). Trial and intertrial durations in autoshaping. *Journal of Experimental Psychology – Animal Behavior Processes*, 3(3), 264–284.

Gibbon, J., & Balsam, P. (1981). Spreading association in time. *Autoshaping and Conditioning Theory*, 219–253.

Gibbon, J., Farrell, L., et al. (1980). Partial-reinforcement in autoshaping with pigeons. *Animal Learning Behavior*, 8(1), 45–59.

Gluck, M. A., & Thompson, R. F. (1987). Modeling the neural substrates of associative learning and memory: a computational approach. *Psychological Review*, 94(2), 176–191.

Gottlieb, D. A. (2004). Acquisition with partial and continuous reinforcement in pigeon autoshaping. *Learning and Behavior*, 32(3), 321–334.

Gottlieb, D. A. (2005). Acquisition with partial and continuous reinforcement in rat magazine approach. *Journal of Experimental Psychology: Animal Behavior Processes*, 31(3), 319–333.

Gottlieb, D. A., & Prince, E. B. (2012). Isolated effects of number of acquisition trials on extinction of rat conditioned approach behavior. *Behavioural Processes*, 90(1), 34–48.

- Gottlieb, D. A., & Rescorla, R. A. (2010). Within-subject effects of number of trials in rat conditioning procedures. *Journal of Experimental Psychology: Animal Behavior Processes*, 36(2), 217–231.
- Harris, J. A. (2011). The acquisition of conditioned responding. *Journal of Experimental Psychology: Animal Behavior Processes*, 37, 151–164.
- Haselgrove, M., Aydin, A., et al. (2004). A partial reinforcement extinction effect despite equal rates of reinforcement during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 30(3), 240–250.
- Hawkins, R. D., Kandel, E. R., et al. (2006). Molecular mechanisms of memory storage in *Aplysia*. *Biological Bulletin*, 210(3), 174–191.
- Helmholtz, H. v. (1852). Messungen über Fortpflanzungsgeschwindigkeit der Reizung in den Nerven. *Archiv für Anatomie, Physiologie, und wissenschaftliche Medizin*, 199–216.
- Hermann, L. (1879). *Handbuch der Physiologie. Band 2. Teil 1*. Leipzig: Vogel.
- Holder, M. D., Bermudez-Rattoni, F., et al. (1988). Taste potentiated noise-illness associations. *Behavioral Neuroscience*, 102(3), 363–370.
- Holland, P. C. (1980). CS-US interval as a determinant of the form of Pavlovian appetitive conditioned responses. *Journal of Experimental Psychology: Animal Behavior Processes*, 6(2), 155–174.
- Holland, P. C. (2000). Trial and intertrial durations in appetitive conditioning in rats. *Animal Learning and Behavior*, 28(2), 121–135.
- Judson, H. (1980). *The eighth day of creation*. New York: Simon & Schuster.
- Kamin, L. J. (1969a). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 276–296). New York: Appleton-Century-Crofts.
- Kamin, L. J. (1969b). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42–64). Halifax: Dalhousie University Press.
- Kaplan, P. S. (1984). Importance of relative temporal parameters in trace autoshaping – from excitation to inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, 10(2), 113–126.
- Kimble, G. A. (1961). *Hilgard and Marquis' conditioning and learning*. NY: Appleton-Century-Crofts.
- Kirkpatrick, K., & Church, R. M. (2000). Stimulus and temporal cues in classical conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, 26(2), 206–219.
- Laje, R., Cheng, K., et al. (2011). Learning of temporal motor patterns: an analysis of continuous versus reset timing. *Frontiers in Integrative Neuroscience*, 5, 61.
- Lattal, K. M. (1999). Trial and intertrial durations in pavlovian conditioning: issues of learning and performance. *Journal of Experimental Psychology: Animal Behavior Processes*, 25(4), 433–450.
- Lustig, C., Matell, M. S., et al. (2005). Not “just” a coincidence: frontal-striatal interactions in working memory and interval timing. *Memory*, 13(3–4), 441–448.
- Matell, M. S., & Meck, W. H. (2004). Cortico-striatal circuits and interval timing: coincidence detection of oscillatory processes. *Brain Research. Cognitive Brain Research*, 21(2), 139–170.
- Molet, M., & Miller, R. R. (2013). Timing: an attribute of associative learning. *Behavioural Processes*, 95.
- Mustaca, A. E., Gabelli, F., et al. (1991). The effects of varying the interreinforcement interval on appetitive contextual conditioning in rats and ring doves. *Animal Learning and Behavior*, 19, 125–138.
- Ohyama, T., & Mauk, M. (2001). Latent acquisition of timed responses in cerebellar cortex. *Journal of Neuroscience*, 21(2), 682–690.
- Pavlov, I. P. (1927). *Conditioned reflexes*. New York: Dover.
- Raby, C. R., Alexis, D. M., et al. (2007). Planning for the future by western scrub-jays. *Nature*, 445(7130), 919–921.
- Rescorla, R. A. (1967). Inhibition of delay in Pavlovian fear conditioning. *Journal of Comparative and Physiological Psychology*, 64(1), 114–120.
- Rescorla, R. A. (1972). Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 6, pp. 1–46). New York: Academic.
- Rieke, F., Warland, D., et al. (1997). *Spikes: exploring the neural code*. Cambridge MA: MIT Press.
- Silva, K. M., & Timberlake, W. (1997). A behavior systems view of conditioned states during long and short CS-US intervals. *Learning and Motivation*, 28(4), 465–490.
- Silva, K. M., & Timberlake, W. (1998). The organization and temporal properties of appetitive behavior in rats. *Animal Learning and Behavior*, 26(2), 182–195.
- Silva, K. M., & Timberlake, W. (2005). A behavior systems view of the organization of multiple responses during a partially or continuously reinforced interfood clock. *Learning and Behavior*, 33(1), 99–110.
- Silver, R., Balsam, P. D., et al. (2011). Food anticipation depends on oscillators and memories in both body and brain. *Physiology and Behavior*, 104(4), 562–571.
- Staddon, J. E., & Zhang, Y. (1991). On the assignment-of-credit problem in operant learning. In Commons, Michael, L., Grossberg, Stephen, et al. (Eds.), *Neural network models of conditioning and action. Quantitative analyses of behavior series* (pp. 279–293) (xx, p. 359).
- Stein, L., Sidman, M., et al. (1958). Some effects of 2 temporal variables on conditioned suppression. *Journal of the Experimental Analysis of Behavior*, 1(1), 153–162.
- Sutton, R. S. (1984). Temporal credit assignment in reinforcement learning. PhD, University of Massachusetts, Amherst.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3, 9–44.
- Wagner, A. R., Logan, F. A., et al. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, 76(2), 171–180.
- Ward, R. D., Gallistel, C. R., et al. (2012). Conditioned [corrected] stimulus informativeness governs conditioned stimulus-unconditioned stimulus associability. *Journal of Experimental Psychology: Animal Behavior Processes*, 38(3), 217–232.
- Ward, R. D., Gallistel, C. R., et al. (2013). It is the information! *Behavioural Processes*, 95, 3–7.
- Williams, B. A. (1981). Invariance in reinforcements to acquisition, with implications for the theory of inhibition. *Behaviour Analysis Letters*, 1, 73–80.