Neurobiology of Learning and Memory 108 (2014) 136–144

Contents lists available at ScienceDirect

Neurobiology of Learning and Memory

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

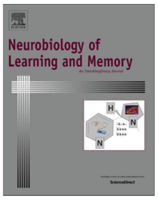

Review

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.

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 pairing 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).

1074-7427/$ - see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.nlm.2013.11.019

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 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 parings 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