

The Replacement of General-Purpose Learning Models with Adaptively Specialized Learning Modules

C.R. Gallistel

Theories of learning are and always have been predominantly associative theories. However, in the study of animal learning, where these theories have historically been most dominant, a different conception is gaining ground. Whereas associative theories have their historical roots in the empiricist philosophy of mind, the alternative conception has its roots in evolutionary biology, more particularly in zoology, that is, in the study of the natural history of animal behavior and of the mechanisms that enable animals to cope with the challenges posed by their habits of life.

Associative theories of learning assume a basic learning mechanism, or, in any event, a modest number of learning mechanisms. These mechanisms are distinguished by their properties--for example, whether or not they depend on temporal pairing--not by the particular kind of problem their special structure enables them to solve. Indeed, people doing neural net modeling, which is currently the most widespread form of associative theorizing, are often at pains to point out that the network has solved a problem in the absence of an initial structure tailored to the solution of that problem (e.g. Becker & Hinton, 1992).

The alternative conceptualization, by contrast, takes for granted that biological mechanisms are hierarchically nested adaptive specializations, each mechanism constituting a particular solution to a particular problem. The foliated structure of the lung reflects its role as the organ of gas exchange, and so does the specialized structure of the tissue that lines it. The structure of the hemoglobin molecule reflects its function as an oxygen carrier. The structure of the rhodopsin molecule reflects its function as a photon-activated enzyme. One cannot use a hemoglobin molecule as the first stage in light transduction and one cannot use a rhodopsin molecule as an oxygen carrier, any more than one can see with an ear or hear with an eye. Adaptive specialization of mechanism is so ubiquitous and so obvious in biology, at every level of analysis, and for every kind of function, that no one thinks it necessary to call attention to it as a general principle about biological mechanisms.

In this light, it is odd but true that most past and contemporary theorizing about learning does not assume that learning mechanisms are adaptively specialized for the solution of particular kinds of problems. Most theorizing assumes that there is a general purpose learning process in the brain, a process adapted only to solving the problem of learning. Needless to say, there is never an attempt to formalize what exactly that problem is. From a biological perspective, this is equivalent to assuming that there is a general purpose sensory organ, which solves the problem of sensing.

In this chapter, I review some of the evidence that whenever learning occurs, it is made possible by an adaptively specialized learning mechanism--a learning module--whose structure is as specific to a particular learning problem as the structure of a sensory organ like the eye or the ear is specific to a particular modality. The review focuses on the differences in computational structure between several distinct learning mechanisms that are important to an animal's ability to find its way about and a different set of

mechanisms that enable it to learn and exploit the temporal structure of its experience. I finish with some observations on the implications of this view for research on the neural mechanisms of learning and memory.

Spatial Learning Mechanisms

Learning by Path Integration

Our first example of a learning module is so far removed from what is commonly regarded as a learning mechanism that before I discuss it I commonly get people to agree that the pretheoretic definition of a learning mechanism is a mechanism by which we acquire knowledge of the world and our place in it. Then, I show them figure X-1, which is the track of a long-legged fast-moving desert ant, *Cataglyphis bicolor*, foraging for and finding a morsel of food on the hot plain of the Tunisian desert. On the outward leg of its journey (solid tracing), it twists and turns this way and that searching for the carcass of an insect that has succumbed to the heat. When it finally finds one, it bites off a chunk, turns, and runs more or less straight toward its nest, a hole 1 mm in diameter, which may be as much as 50 m away. Its ability to orient homeward demonstrably depends on information it acquires during the outward journey. If the ant is deprived of this information by being picked up as it emerges from the nest and transported to an arbitrary point in the vicinity of its nest, it wanders in circles and makes it back to the nest only after a long time, if at all (Wehner & Flatt, 1972). A simple experiment by Wehner and Srinivasan (1981) reveals the nature of the learning mechanism that acquires the requisite positional information. When the ant turns from a food source back toward its nest, it is picked up and transported more than half a kilometer across the desert, here it is released to run across a large grid the experimenter has marked in advance on the desert floor. The grid on the desert enables the experimenter to trace the ant's course on a graph-paper grid at a reduction of 100:1 (figure X-2). Although the ant is now in territory it has never seen, it runs in a direction that lies within a degree or two of the compass direction of its nest from the site where it was picked up, the direction in which it would have run had it not been displaced. It runs in a straight line for a distance slight longer than the distance of its nest from the point where it was picked up, the abruptly stops its straight run and begins a systematic search for the nonexistent nest.



FIGURE X.1. *Track of a foraging ant. The outward (searching) journey is the solid line. It found food at X. Its homeward run is the broken line. Redrawn from Harkness and Maroudas (1985). Used by permission of author and publisher.*

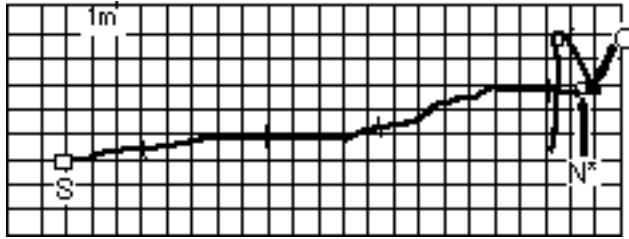


FIGURE X.2. *The homeward run of an ant displaced into unfamiliar territory on which the experimenters had previously marked of a 1 meter gridwork. S indicates the point of release. N* indicates the location of the fictive nest (where the ant's nest would have been had the ant not been displaced). O indicates the point at which the ant terminated its nestward run and entered a search pattern. Redrawn from Wehner & Srinivasan, 1981. Used by permission of author and publisher.*

The ant's ability to run over unfamiliar terrain a course whose direction and distance equal the direction and distance of its nest from the point of its capture implies that its navigation is based on some form of path integration or dead reckoning. Path integration is the integration of the velocity vector with respect to time to obtain the position vector, or a discrete approximation to this computation. The discrete approximation in traditional marine navigation is to record the direction and speed of travel (the velocity) at intervals, multiply each recorded velocity by the interval since the previous recording to get interval-by-interval displacements (e.g., making 5 knots on a northeast course for half an hour puts the ship 2.5 nautical miles northeast of where it was), and sum the successive displacements (changes in position) to get the net change in position. These running sums of the longitudinal and latitudinal displacements are the deduced reckoning of the ship's location (abbreviated "ded. reckoning" whence, "dead reckoning").

For this computation to be possible, the ant's nervous system must have elements capable of preserving the value of a variable over time. The essence of the ongoing summation that underlies path integration is the adding of values that specify the most recent displacement to the values that specify the cumulative prior displacement. To do this, one must be able to hold in memory the value of the sum, and add to that value.

Path integration is a process that computes and stores values that specify a quantifiable objective fact about the world--the ant's direction and distance from its nest. The use made of the values, if any, is determined by the decision process controlling behavior at the moment. The integration of velocity to obtain position goes on throughout the ant's journey, but only when the decision to turn for home is made do the positional values stored in the dead-reckoning integrator lead to the well-oriented straight runs for a predetermined distance shown in figures X.1 and X.2. The position-specifying values in the nervous system are in no sense associative bonds: they were not created by temporal pairing and they do not function as conducting links in the circuitry that links sensory input to behavioral output, nor in circuitry that links the memory of one sensory input to the memory of a different sensory input. Rather, these position-specifying values in the nervous system--the product of the dead reckoning module--are symbols. Their values (magnitudes) specify a fact about the world, which different decision processes may use in different ways. They are no more conducting links than are memory locations in a computer or genes on a chromosome. Like these two examples, they are repositories of information.

The values accumulated in the integrator at important points in the forager's journey may also be stored for later use-- either in the direction of overt behavior or in computations that lead to the storing of other facts about the world. For example, the ant stores the positional coordinates of the location at which it finds the carcass, so that it can return directly there on its next foraging trip. We infer this because when we set up a food source in a fixed location, the ants soon come directly to it from their nest (Wehner and Srinivasan, 1981). These fast-moving ants do not lay an odor trails, so their ability to set a course directly for the food source from the nest implies that they have recorded its position relative to the nest.

The foraging bee gives direct evidence that it has stored the positional coordinates of a food source for later use. When it returns to its hive, it performs a symbolic dance. It runs in a figure-eight path, wagging when it runs the central segment where the two loops join. Other foragers follow behind the dancing bee to learn the location of the food source. The direction of the wagging part of the run relative to the gravitational vertical specifies the direction of the food source relative to the sun (the solar bearing of the source), while the number of waggles specifies the approximate distance of the source from the hive (Frisch, 1967). In this way, the bee communicates to others the position-specifying values computed and saved by its path-integrating mechanism. This is but one of several possible illustrations of the diverse uses to which values stored in memory may be put by different behavioral read-out systems.

If the adaptive specialization of the complex computations that mediate information acquisition is taken for granted, the the only plausible elementary mechanisms that might be common to learning mechanisms in general is a mechanism for storing the computed values of variables and the mechanisms that carry out the primitive elements of all computation (adding, subtracting, multiplying, dividing, ordinating, conditional testing, and so on). Path integration is a notably simple and widespread neural computation that makes straightforward use of these elementary computational operations. It ought therefore to be of central interest to the large community of scientists who seek to discover the cellular-level mechanism that make learning and higher cognitive function possible. It ought in other words to be regarded as a paradigmatic learning process worthy of intensive study at the neurobiological level.

Why then do many people find the idea of path integration as a paradigmatic and notably simple instance of learning perverse? The reasons offered are revelatory. First, it is argued that it is a special-purpose innate mechanism designed to do this and only this. In short, it's an adaptive specialization. This is equivalent to arguing that the eye is not a representative example of a sensory organ because it is designed specifically for vision. It presupposes that there is such a thing as a general purpose learning mechanism, precisely the presupposition that the zoological approach to learning calls into question.

Another reason offered for regarding path integration as unusual is that the ant does not have to experience its outward journey repeatedly before it knows the way. A variant of this argument is that the ant does not get better at running home with repeated practice. These objections reify the basic assumption of associative theory, which is that learning involves the gradual strengthening of something. To be sure, the conditioned response in classical conditioning does not usually appear until there have been several co-occurrences of the conditioned and unconditioned stimuli, but the only apparent justification for making strengthening through repetition part of the definition of learning is the conviction that there is a general purpose learning mechanism and that the classical conditioning paradigm captures its essence. If we think that path integration captures the only essence of learning that is there to be captured, then we are not going to make strengthening by repetition part of the definition of learning. Moreover, as we will

see later, the domain-specific approach to learning mechanisms offers an interpretation of classical conditioning in which the notion of strengthening a conductive connection through repeated pairing of stimuli plays no role.

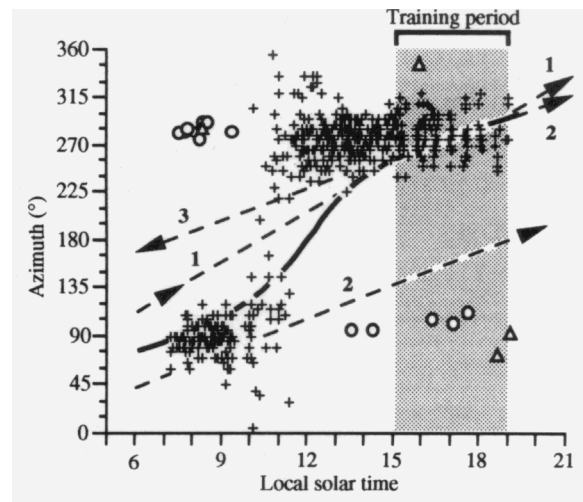
Learning the Solar Ephemeris

Dead reckoning requires a fixed directional reference; the animal must know which way is north at all times. A wide variety of animals, including the ant, derive this information from the sun's azimuth, the point on the horizon above which the sun is positioned at a given moment. This is remarkable, because to use the sun for directional reference, the animal must know the solar ephemeris, the sun's position as a function of the time of the year and the day. Because the sun's azimuth at a given time of day depends strongly on the animal's latitude and the time of year, the solar ephemeris must be learned. A recent experiment by Dyer and Dickinson (1994) sheds new light on the mechanism that is specialized to learn this.

It has long been known that one of the characteristics of the mechanism that learns the solar ephemeris is that it enables animals to judge the position of the sun at times of day when they have never seen--and will never see--the sun. Honey bees at mid-northern latitudes, for example, treat a sun substitute presented at midnight on their internal clock as if it were due north (M. Lindauer, 1957). And bees that have only experienced the sun in the afternoon, when it is in the west, assume that it is nonetheless in the east in the morning (M Lindauer, 1959). The most common explanation of this has been that the bee's learn the average angular velocity of the sun's azimuth (the number of degrees it moves along the horizon per hour) and extrapolate its position by multiplying its average velocity by the time elapsed since they last saw it. This explanation has always been somewhat problematic, because at tropical latitudes the sun's angular velocity is negligible throughout the late afternoon, and bees evolved in the tropics.

In any event, the Dyer and Dickinson experiments show that it is not extrapolation but rather curve-fitting that enables the bees to know the sun's direction at times of day when they have never observed it. Dyer and Dickinson restricted the foraging experience of incubator raised bees to the last 4 hours of daylight, during which the sun was always more or less in the west, with its azimuth moving through less than 20% of the angle it covers between dawn and dusk. The bees foraged from a single source in a fixed location. As already noted, the dance the bees execute when they return to the hive indicates the solar bearing of the source. For example, if the source is due south of the hive (compass bearing 180°) and the sun is due west (compass bearing 270°), then the dance is at an angle of $180^\circ - 270^\circ = -90^\circ$, or 90° to the left of vertical, the number of degrees that the source is counterclockwise from the current position of the sun. When the compass bearing of the source is known, then the observed angle of the dance gives the bee's estimate of the sun's azimuth. On heavily overcast days, when bees cannot see the direction of the sun (Brines & Gould, 1982), returning foragers dance nonetheless. This is possible precisely because they have learned the solar ephemeris. They know where the sun is when they cannot see it. In the early morning of such a day, Dyer and Dickinson released the bees who had seen the sun only in the afternoon. The released bees foraged from the fixed source throughout the day, and the changing angle of their dance indicated their changing estimate of where the sun was. That is, it indicated the solar ephemeris the bees had derived from their limited experience. The results are in Figure X.3

FIGURE X.3 *The compass direction (Azimuth) of the feeding source indicated by individual bee dances, as a function of the time of day, in bees with prior experience only with the afternoon sun (Training period). From Dyer & Dickinson (1994) by permission of the authors and publisher.*



As in Lindauer's classic experiments, bees who had never experienced the sun except in the afternoon, when it is in the west, nonetheless believed it to be in the east in the morning. But as the morning wore on, they did not think that it moved steadily into the south, as they would if they extrapolated the previously observed motion of the sun in any of the ways that have been previously proposed (dashed lines). In the solar ephemeris these bees had derived, the sun stayed in the east until about noon, when it moved abruptly into the west, changing its azimuth by 180° in a very short time. In short, the mechanism for learning the solar ephemeris has built into it what is universally true about the sun, no matter where one is on the earth: it is somewhere in the east in the morning and somewhere in the west in the afternoon. Learning the solar ephemeris is a matter of adjusting the parameters of this a priori function so as to make it fit the observed directions of the sun.

This is reminiscent of the mechanism by which birds learn to sing a song appropriate to their species and region (Marler, 1991), despite the fact that they hear songs from many other species, some closely related, during the period when they are learning their own song. Built into the learning mechanism is a template or selective filter that enables the birds to recognize the kind of song they should be learning, the features that are universal in the song of that species. It is also reminiscent of contemporary theories of language learning, which assume that the learning of a human language is a matter of establishing through observation the parameter values that enable a universal grammar to generate the language spoken by those around you. A universal grammar has built into it what is universally true about human languages.

If problem specific learning mechanisms are required to explain everything from the learning of the solar ephemeris in bees to song learning in birds and language in humans, why should we imagine that there exists a general purpose mechanism in addition to all these problem specific mechanisms? And what structure could it have? It is like trying to imagine what the structure of a general purpose organ might be, the organ that took care of the problems not taken care of by adaptively specialized organs like the liver, the kidney, the heart, and the lungs.

Temporal Learning Mechanisms

Associative learning that has been assumed to serve quite general purposes. An association is a connection between two units of mental or neural activity (two ideas, two neurons, two nodes in a neural net, etc.). The associative connection arises either because the two units have often been active at nearly the same time (the temporal pairing of activation) or through the repeated operation of a feedback mechanism that is activated by errors in the output and adjusts associative strengths to reduce the error. In traditional animal learning theory, the process that forms associations by the first mechanism (temporal pairing of activation) is called the classical or Pavlovian conditioning process, while the process that forms associations through the agency of error-correcting feedback is called the instrumental or operant conditioning process. In neural net modeling, the first process is called an unsupervised learning mechanism (e.g., (Becker & Hinton, 1992)) in contrast to the more common second mechanism, which requires a supervisor or teacher that knows the correct output. The correct output is required for the error-actuated back-propagation algorithm, which plays the role of the reinforcement process, selectively strengthening those connections that lead to correct outputs. In theories of instrumental learning (Hull, 1943), the error correcting feedback comes from the rewards and punishments generated by effective acts.

The Pavlovian and operant experimental paradigms were created to study the principles of association formation. The creators shared with contemporary connectionist modelers the assumption that the general nature of learning was already known; it was associative. The problem was to determine the details of the associative process and how it could explain whatever was observed. The paradigms were not intended to be laboratory analogs of particular learning problems that animals had been observed to solve in the field. This does not mean, of course, that the paradigms do not in fact instantiate problems of a particular kind. In fact they do. From a mathematical perspective, they are examples of problems in multivariate, non-stationary time series analysis.

They are time series problems because what the animals are learning is the temporal dependence of one event on another. One of the most important discoveries to emerge from the modern study of conditioning is that temporal pairing is neither necessary nor sufficient for conditioning, because animals in conditioning experiments learn the contingencies between stimuli, and contingency cannot be reduced to temporal pairing. The contingency between a conditioned stimulus (CS) and an unconditioned stimulus (US) is just as great if the US occurs only when the CS is absent (the explicitly unpaired protocol) as it is if the US occurs only when the CS is present (the standard protocol). Both contingencies are learned (Colwill, 1991; LoLordo & Fairless, 1985; Rescorla, 1969), but in the first, the CS and US are never paired, while in the second, they are always paired. Thus, temporal pairing is not necessary.

They are multivariate time series problems, because there are many different events or time-varying conditions that may or may not predict the time or rate of US occurrence. An important challenge the animal faces is the problem of figuring out what predicts what. This is the multivariate problem. The modern era in the study of conditioning began with the almost simultaneous discovery by three different laboratories of the fact that the learning mechanism operating during conditioning experiments solves the multivariate problem in ingenious and sophisticated ways (Kamin, 1969; Rescorla, 1968; Wagner, Logan, Haberlandt, & Price, 1968). One of the things established by the almost simultaneous discovery of blocking, overshadowing, the effects of background conditioning and of the relative validity of CSs was that temporal pairing is not sufficient.


Conditioning protocols require solutions to the problem of non-stationarity, because the contingencies between CSs and USs commonly change from one phase of the experiment to another. In extinction experiments, for example, the rate of reinforcement (US delivery) is higher when the CS is present, but only during the training phase of the experiment. During the extinction phase, the rate of reinforcement is zero when the CS is present. This is an example of a non-stationary temporal dependency, a contingency that itself changes over time. It has been known since the early days of the study of both classical and operant conditioning that the subjects detect these changes in contingency, and alter their behavior appropriately. The possible existence of a non-stationary contingencies requires a particularly sophisticated kinds of time series analysis.

The fact that conditioning protocols present problems in multivariate, non-stationary time series analysis is irrelevant from the perspective of associative models of the conditioning process, because the associative process is not assumed to be specialized for this or any other particular sort of problem. For that very reason, it has proved difficult to elaborate an associative model of the conditioning process capable of explaining in an internally consistent, mathematically simple and economical way the results from conditioning experiments []. In recent years, models that take as their point of departure the assumption that the results are generated by a learning mechanism specifically tailored to solve this kind of problem have been more successful []. They naturally explain fundamental features of the conditioning data that have challenged associative models for decades.

The problems that associative models confront in attempting to explain the results of classical and operant conditioning experiments are legion [], but very many of them stem from the time-scale invariance of the conditioning process. This means that when one conducts the same experiment at different scales--multiplying or dividing all the time intervals in the protocol by a common factor--one gets the same results. The results are the same (invariant) regardless of the time scale, sometimes over as much as two orders of magnitude variation in the time scale.

The simplest example of the time-scale invariance of conditioned behavior, and the first to be noted (Dews, 1970; J. Gibbon, 1972, 1977), is in the timing of the conditioned response. If, as is commonly the case, reinforcement is given at a fixed interval after the onset of a conditioned stimulus or at a fixed interval after the last reinforcement (e.g., on a fixed interval or FI schedule), then the conditioned response (salivation or key pecking or lever pressing or an eye blink) begins on average at some fixed proportion of the reinforcement latency. Because the ratio between the reinforcement latency and the onset of responding is constant, the time scale does not matter--the scaling factor is common to both the numerator and the denominator of the fraction. The moment at which responding begins varies substantially from trial to trial, but this variability is also proportional to reinforcement latency. Thus, when the distributions of response onsets (and offsets) are normalized, that is, when the time scale is taken out, they are superimposable (Figure X.4). Thus, the noise or variability in the timing of the conditioned response is time-scale invariant (J. Gibbon, 1977; J. Gibbon, Church, & Meck, 1984).


FIGURE X.4. Dew's (1970) data on the normalized rate of responding during the interreward interval on fixed interval schedules ranging from 30 s to 3,000 s, as a function of the normalized time in the interval. The normalized rate is the average rate at elapsed time, t , since the most recent reward divided by the average rate at latency T , just before the next reward. The normalized time is t/T . The responding in any one inter-reward interval is discontinuous. The seemingly smooth acceleration of responding seen here is an averaging artifact. What these data show is that the distribution of response onsets are time scale invariant, because the normalized distributions superimpose. (Reproduced with minor modifications from J. Gibbon, 1977 and used by permission of the author and publisher.)



Explaining the fact that the timing of the onsets and offsets of conditioned responding is proportional to reinforcement latency is one of the challenges that associative models of the conditioning process have not confronted. It is difficult, because it implies that the subjects remember the reinforcement latency. Because associative bonds are assumed to be conducting links whose current strengths depends on many different aspects of the subjects conditioning experience, they cannot specify reinforcement latency. And, indeed, reinforcement latency (the delay of reinforcement) is commonly assumed to be one of the many variables that determine associative strength, rather than a variable whose value is encoded by associative strength.

Contrary to what is generally imagined (e.g., Usherwood, 1993, p. 427), the acquisition process in conditioning is also time scale invariant. Changing the durations of the intervals in a conditioning protocol--for example, the delay of reinforcement--has no affect on the rate of conditioning (the inverse of reinforcements to acquisition) provided one does not change the temporal proportions, how much of one kind of interval there is relative to each other kind. It is commonly supposed that delaying reinforcement retards or prevents the acquisition of conditioned responding, but this is because experiments that seemed to show this (e.g., Coleman & Gormezano, 1971; Schneiderman & Gormezano, 1964) did not change the interval between trials in proportion to the changes in the delay between CS onset and the US. When this interval is increased in proportion to the increase in the delay of reinforcement, delay of reinforcement has no effect on the rate of acquisition (Figure X.5).

FIGURE X.5. Reinforcements to acquisition in the pigeon autoshaping paradigm as a function of delay of reinforcement, with intertrial interval (I) fixed and with intertrial interval varying in proportion to the delay (I/T fixed). From Gallistel & Gibbon,, 1998a, under review.



The time-scale invariance of the acquisition process is also seen in the fact that partial reinforcement does not affect either the rate of acquisition or the rate of extinction. The number of reinforcements required for acquisition is not affected by inserting as many as 9 unreinforced trials for every one reinforced trial during training (more or less flat dashed lines in Figure X.6B). Moreover, the proportion of reinforced trials during training has no effect on how many reinforcements must be omitted to induce the subject to stop responding during extinction (more or less flat dashed lines in Figure X.6C).

The failure of partial reinforcement to affect the rates of acquisition and extinction is profoundly puzzling from an associative perspective, because the effects of unreinforced trials are assumed to antagonize the effects of reinforced trials. Intermingling unreinforced trials with the reinforced trials during acquisition should augment the number of reinforced trials required for acquisition. It should also make the net excitatory strength of the associative bonds after any given number of reinforced trials weaker the more unreinforced trials have been intermingled, so that it takes fewer unreinforced trials to produce extinction. But neither result is in fact seen, because partial reinforcement does not change the relative amounts of exposure to the CS and to the background. It increases the amount of exposure to the CS per reinforcement by the thinning factor. That is, there is twice as much exposure to the CS per reinforcement with a 1:2 schedule of partial reinforcement. But it also increases the amount of exposure to the background alone by the same factor. (That is, there is also twice as much exposure to the background alone per CS reinforcement.) Because conditioning depends only on these temporal proportions (time-scale invariance), and these proportions are not altered, partial reinforcement affects neither the rate of acquisition nor the rate of subsequent extinction.

Explaining the time-scale invariance of conditioning is difficult for associative models for three interrelated reasons: (1) They usually assume that temporal pairing is important. (2) They invariably assume that it is differences in probability of reinforcement that are important. (3) They almost invariably divide continuous time into discrete trials. Each of these assumptions directly violates time-scale invariance. The notion of temporal pairing assumes that there is some time window within which the CS and US must both occur if they are to be regarded as temporally paired. The width of this assumed window imposes a time scale. Probabilities of reinforcement and differences in probabilities of reinforcement cannot be defined except by assuming some discrete finite interval of time within which the probability is measured. That is, they cannot be defined without assuming that continuous time is divided into discrete trials, because the probability of any event goes to zero (as do differences in probability) as the intervals within which probabilities are measured become arbitrarily small. Thus, the values of all probabilities depend on the durations of the trials. The durations of the assumed trials impose a time scale.

If, however, we assume that the mechanism that underlies conditioning is specialized to compute the objective effects of various stimulus conditions on the relative rates of reinforcement and the statistical objective limits on the certainty with which these effects can be known given a limited amount of experience, then conditioning should be time scale invariant, because changing time scales affects neither the relative rates of reinforcement nor the limits on the certainty with which they can be known after any given number of reinforcements have been delivered.




Figure X.6. A. The parameters of a simple conditioning protocol (T, I, C & S). Dots indicate reinforcements which are delivered at CS termination (on reinforced trials). B. The effect of the partial reinforcement schedule (x axis) and the C/T ratio (between curves) on trials to acquisition (solid lines) and reinforcements to acquisition (dashed lines). C. The effect of the partial reinforcement schedule in force during training on unreinforced trials to extinction (solid lines) and on the number of omitted reinforcements to extinction (dashed line). An omitted reinforcement is a reinforcement that would have been delivered by the training schedule. (Reproduced from Gallistel & Gibbon, 1998 by permission of author and publisher.)

The discussion of the mechanisms underlying matching behavior, which has been intensively studied by both zoologists interested in foraging behavior and operant psychologists interested in choice processes, is only one indication of the power of the timing models when applied to data from classical and instrumental conditioning experiments. For example, these models give a more straightforward explanation of the effects of background conditioning (Rescorla, 1968), which are crucial in modern analyses of conditioning. When USs are given in the presence of the background alone (that is to the animal while it is in the apparatus but when there is no explicit CS), the animal "associates" these USs with the background and this association may block the learning that would otherwise occur when the US is paired with an explicit CS. The associative analysis of this important phenomenon requires an internal trial clock (Rescorla & Wagner, 1972), for which there is no evidence. The interval-timing analysis (Gallistel, 1990; J. Gibbon & Balsam, 1981) simply assumes that the animal learns the inter US interval to be expected in the apparatus itself and does not exhibit a conditioned response to another CS if the inter US interval in the presence of that CS (and, of course,

the background) is no different from the inter US interval predicted by the apparatus (background).

In associative theories of classical conditioning, it is assumed that the temporal interval between CS onset and US onset determines the rate of conditioning, but it has been shown that there is no effect of the CS-US interval on the rate of conditioning if the US-US interval is adjusted in proportion to the change in CS-US interval (J. Gibbon, Baldock, Locurto, Gold, & Terrace, 1977). Put another way, the rate of conditioning is invariant under scalar transformation of the temporal intervals in the training protocol. Explaining this scalar invariance is a major challenge to associative models, but scalar invariance in the rate of conditioning is a straightforward prediction of the timing models (Gallistel, 1990, 1992; J. Gibbon & Balsam, 1981).

Making quantitative predictions about the effects of partial reinforcement on the rate of conditioning is another major challenge to associative models. Timing models, on the other hand, predict that partial reinforcement will have a scalar effect on the rate of conditioning. A scalar effect of partial reinforcement on the rate of conditioning is what is observed, at least in autoshaping paradigms (Gibbon, J., Farrell, Locurto, Duncan, & Terrace, 1980). That is, when the US occurs during only one in every 10 occurrences of the CS, it takes 10 times as many CS occurrences to produce the same strength of conditioned responding.

Partial reinforcement makes conditioning slower, but it also makes extinction slower. Explaining the retarding effect of partial reinforcement on the rate of extinction is a problem that has long bedeviled associative models. Why should the history of conditioning determine how many non-reinforcements it takes to weaken the association a given amount? On the other hand, timing models predict that the partial reinforcement will increase the trials to extinction by a factor proportionate to the thinning of the reinforcement schedule (Gallistel, 1992), which is the empirical results, at least in autoshaping experiments (Gibbon, J. et al., 1980).

Finally, timing models predict that increasing the intertrial interval, which greatly increases the rate of initial conditioning, will have no effect on the rate of extinction. Again, this is the empirical result (Gibbon, J. et al., 1980).

In short, models that assume adaptively specialized rather than general purpose learning mechanisms are now being applied to the prediction and explanation of well established findings in the classical and instrumental conditioning literature. This suggests that learning mechanisms, like other biological mechanisms, will invariably exhibit adaptive specialization. We should no more expect to find a general purpose learning mechanism than we should expect to find a general purpose sensory organ.

References

- Becker, S., & Hinton, G. E. (1992). Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature*, 355(Jan 9), 161-163.
- Brines, M. L., & Gould, J. L. (1982). Skylight polarization patterns and animal orientation. *Journal of Experimental Biology*, 96, 69-91.

- Coleman, S. R., & Gormezano, I. (1971). Classical conditioning of the rabbit's (*Oryctolagus cuniculus*) nictitating membrane response under symmetrical CS-US interval shifts. *Journal of Comparative and Physiological Psychology*, 77, 447-455.
- Colwill, R. M. (1991). Negative discriminative stimuli provide information about the identity of omitted response-contingent outcomes. *19*, 326-336.
- Dews, P. B. (1970). The theory of fixed-interval responding. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 43-61). New York: Appleton-Century-Crofts.
- Dyer, F. C., & Dickinson, J. A. (1994). Development of sun compensation by honeybees: How partially experienced bees estimate the sun's course. *Proceedings of the National Academy of Sciences, USA*, 91, 4471-4474.
- Frisch, K. v. (1967). *The dance-language and orientation of bees*. Cambridge, MA: Harvard Univ. Press.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press.
- Gallistel, C. R. (1992). Classical conditioning as an adaptive specialization: A computational model. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 28, pp. 35-67). New York: Academic Press.
- 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. (1972). Timing and discrimination of shock density in avoidance. *Psychological Review*, 79, 68-92.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's Law in animal timing. *Psychological Review*, 84, 279-335.
- Gibbon, J., Baldock, M. D., Locurto, C. M., Gold, L., & Terrace, H. S. (1977). Trial and intertrial durations in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, 3, 264-284.
- Gibbon, J., & Balsam, P. (1981). Spreading associations in time. In C. M. Locurto, H. S. Terrace & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219-253). New York: Academic.
- Gibbon, J., Church, R. M., & Meck, W. H. (1984). Scalar timing in memory. In J. Gibbon & L. Allan (Eds.), *Timing and time perception* (Vol. 423, pp. 52-77). New York: New York Academy of Sciences.
- Hull, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.

- Kamin, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 42-64). Halifax: Dalhousie University Press.
- Lindauer, M. (1957). Sonnenorientierung der Bienen unter der Aequatorsonne und zur Nachtzeit. *Naturwissenschaften*, *44*, 1-6.
- Lindauer, M. (1959). Angeborene und erlernte Komponenten in der Sonnenorientierung der Bienen. *Zeitschrift für vergleichende Physiologie*, *42*, 43-63.
- LoLordo, V. M., & Fairless, J. L. (1985). Pavlovian conditioned inhibition: The literature since 1969. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Marler, P. (1991). The instinct to learn. In S. Carey & R. Gelman (Eds.), *The epigenesis of mind* (pp. 37-66). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*(1), 1-5.
- Rescorla, R. A. (1969). Pavlovian conditioned inhibition. *Psychological Bulletin*, *72*, 77-94.
- 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.
- Schneiderman, N., & Gormezano, I. (1964). Conditioning of the nictitating membrane of the rabbit as a function of CS-US interval. *Journal of Comparative and Physiological Psychology*, *57*, 188-195.
- Usherwood, P. N. R. (1993). Memories are made of this. *Trends in Neurosciences*, *16*(11), 427-429.
- Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, *76*(2), 171-180.
- Wehner, R., & Flatt, I. (1972). The visual orientation of desert ants, *Cataglyphis bicolor*, by means of territorial cues. In R. Wehner (Ed.), *Information processing in the visual system of arthropods* (pp. 295-302). New York: Springer.
- Wehner, R., & Srinivasan, M. V. (1981). Searching behavior of desert ants, genus *Cataglyphis* (Formicidae, Hymenoptera). *Journal of Comparative Physiology*, *142*, 315-338.