

The Principle of Adaptive Specialization as It Applies to Learning and Memory

C. R. Gallistel

Rutgers University

New Brunswick, NJ

Biological mechanisms are adapted to the exigencies of the functions they serve. The function of memory is to carry information forward in time. The function of learning is to extract from experience properties of the environment likely to be useful in the determination of future behavior. These different functions lead to different manifestations of the biological principle of adaptive specialization.

Adaptive Specialization in the Memory Mechanism

We do not know what the neurobiological mechanism of memory is, so we cannot say how it is adapted to its function. We can, however, look at other mechanisms that serve the same function to see what adaptations they suggest we may find in a mechanism with this function. One such mechanism is computer memory; another is DNA, the molecular mechanism for carrying hereditary information from one generation to the next. Both of these “memory” mechanisms suggest two exigencies and two principles: The exigencies are thermodynamic stability and high density. The principles are that information is information and that information conveyance requires a code.

Thermodynamic Stability

A memory mechanism requires a physical change that endures for as long as the information it encodes is to be preserved. Volatile changes decay in a fraction of the interval over which the information is to be preserved. An example, of such volatile memory in some computers is capacitive memory, where the information is encoded by the charge pattern on a bank of capacitors. Capacitive memory banks are volatile because the charge on a capacitor decays rapidly; when the charge pattern is no longer readable, the information has been lost. A volatile memory makes sense when it is a buffer where information is ordinarily overwritten by new information within an interval shorter than the decay time. When the same information must be held for longer than the decay time, a volatile memory must be continually refreshed, which uses energy and generates heat. For this reason, computer memories whose function is to preserve large amounts of information for long periods of time are non-volatile: they rely on thermodynamically stable physical changes.

Volatility is a consideration in assessing the plausibility of a memory mechanism often employed in connectionist or neural net modeling. A conventional computer has one or more memory mechanisms, components specifically devoted to carrying information forward in time. A neural net has no component devoted by design to this

function. This lack reflects the absence of an established memory mechanism in our current understanding of neurobiology and the bottom-up commitment in neural net modeling, the attempt to build computational models on a neurobiological foundation.

Whether neural net modeling is wise to embrace a constraint that may someday be seen as having arisen from simple ignorance remains to be seen, because dispensing with a memory mechanism poses computational challenges. On the rare occasions when the environment provides all the information needed to perform a computation (or determine an output) at one time—in one input vector—the absence of a memory poses no problem. In many situations (arguably in almost all), however, the information is given piecemeal, at unpredictable intervals. Then, the absence of a memory poses a problem. A common solution is to have the activity engendered by the information that arrives first circulate in reverberating loops until the rest of the information required to determine the output arrives. A reverberating loop is a highly volatile mechanism; in the absence of the energy required to maintain continuous signal transmission, the signal dissipates in a matter of milliseconds. Reverberating loops have a long history in neurobiological speculation about memory (Hebb, 1949), but their energetic cost is seldom considered. That cost weighs against the hypothesis that large amounts of information are carried forward over long intervals by this mechanism.

Turning our attention to DNA, the only currently known mechanism for carrying large amounts of information forward in time within biological systems, we find again that the mechanism is not volatile. Absent assaults by active agents of degradation like bacteria, the information in organismic DNA is preserved for hundreds or even thousands of years after the death of the organism, which is to say, after the cessation of energy input. The thermodynamic stability of this mechanism is the more remarkable in that so many other aspects of cellular activity and function are highly dynamic. Living is a highly dynamic and energetically costly business, but the mechanism for carrying forward the information necessary to build a living system is neither. That is the key to its successful functioning.

In sum, the known mechanisms for carrying large amounts of information forward through long intervals of time in both artificial and biological systems suggest that thermodynamic stability is a strong exigency for such mechanisms. Once the information is registered, its continued preservation should require little or no metabolic expenditure.

High Density

The spectacular and continuing increase in the density of information storage has been a salient aspect of the technological history of the last five decades. The powerful, small modern computer is made possible by it. There is reason to expect that this increase will continue for another several decades, because the information density in the best modern chips and on the best modern drives is still many orders of magnitude lower than the density of storage achieved in DNA. DNA stores information at the sub-molecular level. At each base-pair position in the helix, there are four possible base pairs, so each position carries 2 bits of information, and the 10 base-pairs in a 34 angstrom long turn of a helix with a 10 angstrom radius carry 20 bits, which works out to about 2 bits per cubic

nanometer (2×10^{18} bits per cubic millimeter). This is within two or three orders of magnitude of what is thought to be the limit on information density that is theoretically possible at ordinary temperatures and pressures (Frank, in press 2002)¹. and about 10 orders of magnitude greater than the density achieved in the highest density contemporary computer memories.

Because of the enduring influence of associative theories of mind on neurobiological theorizing, it is widely hypothesized that the mechanism of memory in neural tissue is a change in synaptic conductance (Kandel, Schwartz, & Jessell, 2000). How much information can be stored in a conductance change at a single synapse depends somewhat on the coding scheme, but it does not appear that it could be more than a few bits. Synaptic densities in neocortical tissue can be on the order of one billion (10^9) synapses per cubic millimeter (Braitenberg & Schüz, 1991). Thus, if every synapse in a cubic millimeter of cortex were devoted to memory storage, the density of information storage would be 8 or 9 orders of magnitude less than is achieved in DNA.

This is an extremely conservative calculation. By which I mean that it is highly biased so as not to make the hypothesis of synaptic storage look too bad. It is widely supposed in contemporary associative theorizing that different states of the world are represented by different stable activity patterns (called attractors) in networks composed of large numbers of interconnected neurons. Information about the world is stored in the patterns of synaptic weights that cause such a network to settle into the activity pattern appropriate to a given input, when it is given that input. A single pattern of weights can support more than one stable state, because the stable state into which the network settles is a function of its connectivity (synaptic weights) and the input pattern it is given. However, there is a calculable limit on the number of different stable states that can be sustained in a network of a given size. The limiting number is a function, primarily, of the number of elements (neurons) in the network and secondarily of the richness of their connectivity. Rolls and Treves. (1998) calculate that the number of different possible attractor states in the rat hippocampus is about 12,000. This means that the rat hippocampus, which occupies a volume of many cubic millimeters, can store less than 16 bits of information, less than can be put into a single turn of the double helix.

A single 32 bit word address in a contemporary computer can represent $2^{32} =$ more than 4 billion different states of the world; whereas, according to calculations of Rolls and Treves, a neural attractor network the size of the rat hippocampus can represent only 12,000. If information in the brain is carried forward by the patterns of connectivity in neural networks like those envisioned by Rolls and Treves and many others, then one would like to understand why the brain uses a mechanism for this critical function that has such spectacularly bad density characteristics.

¹ Current estimates of the density achievable at ordinary temperatures and pressures are about 1 bit per cubic angstrom (Frank, 2000). Densities currently achieved are several orders of magnitude less than this.

In sum, a second exigency in the design of memory mechanisms is the maximization of information density. This gives reason to wonder whether the currently widespread assumption that the memory mechanism must involve a change in synaptic conductance is well founded. If so, then it would seem that the level of information density achieved is many, many orders of magnitude less than what is known to be possible in biological tissue and is currently achieved in manmade devices like, for example, the flash memory in a digital camera.

Information is Information

An important principle in modern computing and communication is that different kinds of information are equivalent and interchangeable when it comes to storage and conveyance; a mechanism suited to store or convey one kind of information is equally well suited to store or convey any other kind. Distinctions between kinds of information are crucial when it comes to mechanisms for encoding and decoding and when it comes to algorithms (computation). For example, in a digital video camera, there is a microphone for encoding the sound stream, whereas there is a 2-dimensional array of light sensitive diodes behind a focusing lens for encoding the luminous flux. And, there are different algorithms for compressing the first-order signals before placing the resulting bit patterns in memory, because the statistics of sound streams and images are different. There are corresponding differences on the decoding side, where the bits stored in memory—and perhaps transmitted over the internet—are converted back into image and sound. But at the level of storage and transmission, a bit is a bit. For the mechanisms that operate at this level, the origin and fate of the information they carry is irrelevant.

One manifestation of this principle in a computer is that the same memory mechanisms are used to store both program code (information specifying the algorithm) and data code (information on which the algorithm operates).

The same is true for biological storage and signaling mechanisms. Very different kinds of information about organismic, cellular and molecular structures and how to build them are carried from generation to generation by the sequence of base pairs in DNA molecules. Some sequences specify protein structure, but many others (the control elements) specify when and where a given protein is to be made.

The function of the action potential is to convey information rapidly over long distances. Here again, only one such mechanism is required. The mechanism that carries information about the sound stream the length of the auditory nerve is the same mechanism that carries information about the visual image the length of the optic tract, and information about mechanical deformation of the skin of the big toe, the length of the body.

Even intracellular and intercellular chemical signaling pathways, where one might expect chemical necessity to operate, are so constructed that in fact none does. The ubiquitous G-protein coupling mechanisms that operate where extracellular chemical signals are converted to intracellular chemical signals remove all chemical connection between the external and internal signals. As a result, within the same basic framework,

any extracellular signal can be made to affect any intracellular cascade. The same principle applies within the intracellular cascades. It has been argued that this principle is a key to understanding how complex organic structures have been able to evolve (Carroll, Grenier, & Weatherbee, 2001).

The principle that information is information gives reason to think that the mechanism for carrying information forward in time within the nervous system may be universal, just as is the mechanism for carrying information rapidly from one place to another. For those interested in the physical basis of memory in brains, the question whether the molecular mechanism is or is not universal is of the first importance. For the sake of scientific progress, one hopes that it is. The information-is-information principle gives some grounds for such a hope.

Coding

The transmission and storage of information are basic elements of information processing. Wherever we have to do with information processing, we have to do with codes. Central to a full understanding of an information storage or conveyance mechanism is an understanding of the code it uses. This was immediately appreciated by leading molecular biologists after the discovery of the structure of DNA. Given the structure of the molecule and the knowledge that it carried the information specifying many different protein sequences, it seemed beyond reasonable doubt that the information about how to make a protein was carried by the sequence of base pairs. But what was the code? How long a sequence of base pairs was needed to specify one amino acid (the building block of a protein)? And what was it about such a sequence that determined which amino acid it specified? And how did the decoding mechanism determine where a coding sequence began and ended? Answering these questions about the code became central to the progress of molecular biology (Judson, 1980).

In neurobiology, the situation is different. Du Bois-Reymond discovered the physical mechanism by which information is carried along nerves in 1848 (du Bois-Reymond, 1848). It was understood at the cellular level by the 1950s (A.L. Hodgkin & A.F. Huxley, 1952a, 1952b; A. L. Hodgkin & A. F. Huxley, 1952; A.L. Hodgkin & A.F. Huxley, 1952c; Hodgkin, Huxley, & Katz, 1952), and the crucial components of the mechanism, the voltage sensitive ion gates, are now well on their way to being understood at the molecular level (Jiang et al., 2002a, 2002b). However, 150 years after the discovery of the transmission mechanism, we have only recently begun to make progress in understanding the code, what it is about a sequence of action potentials that carries the information (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997). Although the nervous system is clearly an information processing system, we have hardly taken the first steps toward understanding how it realizes elementary information processing functions. We do not know the mechanism by which it stores information and, although we have known for 150 years the principal mechanism by which it transmits information over long distances, we are only now beginning to understand the code involved.

The coding principle is important to keep in mind when considering suggested memory mechanisms. One might reasonably ask that they suggest a means by which information may be encoded, since the essential function of a memory mechanism is the conveyance of information. It is widely suggested that Hebbian synapses are the (or an important) mechanism of memory (e.g., Kandel, Schwartz & Jessel, 2000). An objection to this proposal is that Hebbian synapses, as traditionally conceived, lack a fundamental property of any code, intelligibility. A code must code for something. Hebbian synapses, as traditionally conceived, do not and cannot code for anything, because they are the physiological embodiments of associative bonds.

Historically, the strength of an associative bond does not and cannot code for an objectively specifiable property of experience, because it is a one-dimensional consequence of many different experiential factors. It is generally taken to be determined by—at a minimum—the number of trials on which two signals have been temporally paired, how closely paired they were, and how strong the signals were. Mathematically speaking, the function that maps from elementary experience to the strength of an associative bond (or Hebbian synapse) is a many-to-one function. This means that for any one value of the function (any one state of the synapse), there are an infinite number of different combinations of values for the input (numbers of trials, temporal separations, and signal strengths) that could have produced that state. Thus, the state itself, the synaptic conductance, is unintelligible; it cannot tell us anything about the experiences that produced it.

The unintelligibility of the associative bond—the fact that it cannot represent something about the animal's experience—is why associative theories of learning have traditionally been anti-representational (Hull, 1930; Skinner, 1950; Smolensky, 1986), that is, non-symbolic theories. The general idea in an associative theory is that experience rewires the nervous system so as to make it generate better adapted behavior but not so as to encode what it is about the experienced world that makes that behavior adaptive.

This is not to say that changes in synaptic conductance could not be the basis of memory. They could. But to make this possible, the causal conditions for the production of the changes would have to be such as to make the resulting changes correspond poorly to the traditional conception of an associative bond. It was that conception for which Hebb (1949) hypothesized the existence of what have come to be called Hebbian synapses. Thus, the situation here is different from the situation regarding the coding scheme in nerve transmission. There, the difficulty is not in specifying how information might be encoded by the stream of action potentials; the difficulty is in determining how in fact it is encoded. With Hebbian synapses, the difficulty is in specifying how information could in principle be encoded by such a mechanism. What constraints must there be on the encoding process? And how would such a code constrain possible decoding mechanisms?

Adaptive Specialization in Learning Processes

Learning processes are those by which brains acquire knowledge of the animal's environment. They are computational processes because useful knowledge is not directly

expressed in low level sensory signals; it must be extracted from those signals by computational processes. Different kinds of knowledge require different kinds of computations. The structure of these computations reflects the structure of the problem to be solved, just as the structure of a sensory organ reflects the structure of the physical problem that is to be solved (sound gathering, light focusing, and so on).

The associative approach to learning has little use for a concept of adaptive specialization, because it assumes that a single learning process—association formation—mediates all or most learning ((Domjan, 1998, p. 17ff; Hawkins & Kandel, 1984; Miller & Escobar, 2002)). For that to be the case, there would have to be core structural features that constitute the essence of all knowledge acquisition problems. No one, however, has ventured to describe the computational principles common to diverse knowledge acquisition processes because, from an associative perspective, learning is not a matter of knowledge acquisition; it is a matter a plastic nervous system molded by experience to generate more adaptive behavior--but without coding the properties of the world to which the behavior it generates is adapted.

There is now, however, little question that learning does involve knowledge acquisition—the extraction of information from experience. Animals do learn distances and directions in their environment, the times of day at which events happen, the durations of intervals, the numbers of events, and the rates of events (Gallistel, 1990; Gallistel & Gibbon, 2000). Thus, it is no longer possible to ignore the issue of adaptive specialization, because it does not appear to be reasonable to assume that there can be a general purpose knowledge-extraction algorithm, equally suited to extract any kind of information from any kind of input². The following examples illustrate the general principle that it takes different computations to extract different kinds of information from different data. It is these computational principles that act via natural selection over evolutionary time to produce adaptively specialized learning mechanisms.

Learning Current Location

Figure 1 shows a tracing made by Harkness and Maroudas (1985) of the track of a foraging ant of the species *Cataglyphis bicolor* found in the hot and featureless Tunisian desert. This species does not lay odor trails. The tortuous solid line represents the outward journey when the forager was searching hither and yon for carrion; the straight dashed line represents the journey back to the nest after it found something at the location marked with an X. The directness of the return track implies that the ant knew where it was during its wanderings; there was information in its brain specifying the direction back to its nest. This information about its direction from its nest is contingent on the details of its twists and turns on this particular outing, so it cannot have been inherited. It must have been acquired from sensory and efference-copy signals generated as the ant

² For a formal argument why there cannot in principle be general-purpose learning algorithms, see Wolpert, D. H., & Macready, W. G. (1997).

wandered, neural signals that indicated from moment to moment the direction and speed of its progress.




Figure 1. Tracing of a foraging journey by an ant of the species Cataglyphis biolcolor. Solid portion = outward journey (search); dashed portion = return to nest after finding food at point X. (Reproduced from Harkness & Maroudas, 1985, by permission of the authors and publisher.)

Experiments in which foragers of this species were displaced into unfamiliar territory at the start of their homeward journey show that these foragers will run the correct compass course for approximately the correct distance even when they are traversing territory they have never seen (Wehner & Srinivasan, 1981). This implies that the ant's moment-to-moment knowledge of its ever-changing location derives from a process called path integration or dead reckoning. The essence of dead reckoning is the integration of velocity with respect to time. Equivalently, it is the keeping of a running sum of successive small displacements. These displacements are vector quantities, so each displacement must be resolved into components (for example, latitudinal, or north-south, change and longitudinal, or east-west, change). The vector sum at any moment, which is itself a vector, is the animal's net displacement, that is, its change in position, the difference between (for example) its latitude and longitude when it started and its current latitude and longitude.

The circuitry for performing this computation, whatever it may be, is an example of problem-specific knowledge acquisition mechanism. The nervous system extracts a particular kind of information about the animal's relation to its environment—its location within that environment—by a computation that directly reflects principles that govern the world the animal lives in. Position is the integral of velocity with respect to time, quite independent of any computation that the animal may make. This enduring environmental fact no doubt explains the emergence of mechanisms within the brains of animals for integrating velocity signals to obtain position signals, in the same sense as the refractive properties of light explain the evolutionary development of the lenses found in eyes. The path-integration computation (partially) solves this problem and only this problem. No one would suggest, for example, that it is in any way suited to obtain information about the grammar of a locally spoken human language.

Dead reckoning to obtain knowledge of position also illustrates the fundamental role that memory plays in learning. The essence of the computation is the keeping of a running sum. A running sum requires a memory to carry forward in time the previous value of the sum, so that the latest displacement may be added to it.

Learning the Solar Ephemeris

Dead reckoning requires knowledge of the compass direction of one's movement (direction in a framework defined by the earth). There is no point in adding the latest displacement to the sum of previous displacements unless all the displacements have been measured in a frame of reference with a fixed orientation with respect to the earth. Maintaining a fixed frame of reference for estimating one's displacement with respect to the earth is not simple, because, in most environments, where only nearby landmarks can be perceived, their compass direction changes as the animal moves--the principle of motion parallax. Thus, any one feature of the terrain cannot be used as a point of directional reference. The farther away a landmark is, the smaller the parallax change in its direction for a given displacement. Thus, it is highly desirable to use as a point of directional reference something that is a very long ways away, something like the sun or the stars. And, in fact, it has been known for nearly a century that a variety of diurnal animals, including ants and bees, use the sun as the point of directional reference when navigating.

The sun is effectively a point at infinity; hence, it has no parallax motion. Unfortunately, however, it does not have a fixed compass direction. Its compass direction changes as the earth rotates. It can change by more 40 degrees in a single hour. Moreover, the compass direction of the sun at any one time depends in a complicated way on the subject's latitude and on the season of the year. Thus, to use the sun as a point of directional reference, the navigator must know the local solar ephemeris, the sun's compass direction at that location on the earth's surface as a function of the time of day at that season of the year. And, of course, the navigator must also know the time of day.

Knowledge of the time of day is provided by an inborn clock with a period of about 24 hours (Turek, 1998). The clock has a learning mechanism, called the entrainment mechanism, which uses the transitions in overall level of illumination at dawn and dusk to synchronize its oscillations with the day-night cycle. Note that the genes specify the fact about the world that is everywhere and always the same, namely, the period of the day-night cycle, while the experience provides the information that is contingent on one's location, namely the phase of the cycle. To function properly the circadian oscillation in the nervous system must be in phase with the local day-night cycle, and so there is a mechanism that extracts that information from the experience of changes in ambient illumination. This entrainment mechanism is not a general purpose learning mechanism. It extracts this particular kind of information from this particular sensory input and passes it to this particular component of the nervous system (the clock component).

Because the local solar ephemeris, like the local phase of the day-night cycle, depends on where the observer is (and on the season), it is unlikely to be genetically

given. However, it has one component that is everywhere and always true, namely, that the sun is somewhere in the east all morning long and somewhere in the west all afternoon. It turns out that this knowledge, like the period of the day-night cycle, is genetically specified. It can be made behaviorally manifest in the absence of the relevant experience (experience indicating the different directions of the sun at different times of day). These experimental findings suggest that the mechanism for learning the solar ephemeris is a parameter-estimation mechanism; the bees learn from observing the sun's direction relative to the local terrain the appropriate parameters for a genetically specified universal ephemeris function (Dyer & Dickinson, 1996). When the values of its parameters have been correctly set by local observations of the sun's direction at a few widely differing times of day, the function correctly predicts the local direction of the sun as a function of the time of day. But in the absence of observations sufficient to correctly establish the values of its parameters, the built in function, using default parameter values, extrapolates directions for the sun at all times of day, given information about its direction at one time of day.

Dyer and Dickinson (1994) raised bees in an incubator, so they never saw the sun. Then, they allowed them to forage in the latter part of the afternoon, when the sun is more or less due west. They foraged from a feeding station due west of the hive, which means that they flew toward the sun on the way from the hive to the station and away from it on the way back to the hive. In the hive, they danced so as to indicate to other foragers the solar bearing of the food source (its direction relative to the sun). The sun is not visible within the hive, but the gravitational vertical takes the place of the sun in the symbolism of the bees' dance. The dance has the form of a figure 8. The angle relative to the vertical of the middle portion of the dance, where the two loops of the 8 are joined, specifies the solar bearing of the food site. Thus, on these afternoons, the bees ran straight up (symbolically, straight toward the sun) during the critical middle portion of the 8.

The bees were allowed to forage only in the late afternoon, until a day dawned with heavy overcast, which blocked their view of the sun, and which promised to last all day. Now, the bees were allowed to forage for the first time in the morning. The experimenters observed their dance when they returned to the hive from the feeding station. Because the experimenter's knew the compass direction of the feeding station (due west), and because the dance indicates its solar bearing, its direction relative to what the bee takes to be the direction of the sun, the experimenters could infer from the direction of the dance within the hive the direction of the sun specified by the solar ephemeris function in the brain of the dancing bee. In the early morning, when they first began to forage, the bees ran straight down (symbolically straight away from the sun) indicating, more or less correctly, that in order to get to the feedings site it was necessary to fly straight away from the sun.

These bees had never seen the sun in the east and could not see it there on the morning in question. They had never reached the feeding station by flying away from the sun; yet that is what they were in effect telling the other bees to do. That the sun was in the specified direction relative to the local terrain (in the east) was a genetically based inference drawn from the information extracted from their previous afternoon foraging

experiences, during which the sun was in the west. Thus, the direction of the sun at 7:00 in the morning was not information that had been provided by their prior experience of the sun in the afternoon, although it depended critically on that prior experience. Without having seen the direction of the sun (relative to the terrain around the hive) at a known time, and without having visited the feeding site, the bees could have known neither the compass orientation of the local terrain (how it was oriented with respect to the axis of the earth's rotation, which, together with the period of rotation, determines the local solar ephemeris) nor the direction of the feeding site from the hive. On the other hand, without prior knowledge of the universal characteristics of the solar ephemeris, the bees' experience of the sun in the afternoon did not specify how to fly relative to the sun in the morning. In particular, it did not specify that it was necessary to fly away from the sun in the morning even though one flew into the sun in the afternoon. Thus, the bees' estimate of the direction of the sun at 7:00 in the morning depended jointly on information extracted from experience and on an inference (or extrapolation) based on information about the universal properties of the solar ephemeris built into the mechanism that learns the solar ephemeris. It also depended on their inborn clock, which provided the information about the time of day.

As the morning wore on, the bees continued to indicate in their dance that the solar bearing of the feeding station was 180° away from the sun, although in fact, of course, the sun was moving into the south. (The experiments were conducted in the northern hemisphere.) Had they been able to see the sun, that is, had they more accurate knowledge of the local solar ephemeris, they would have danced more and more horizontally, as was known from observation of bees that had seen the sun throughout the day. Somewhat before noon, the directions of the dances became widely dispersed. Some danced as if the sun were in the south, as it would be if they were in the northern hemisphere; while others danced as if it were in the north, as it would be if they were in the southern hemisphere. This period of (justifiable) confusion persisted less than an hour before all the bees had reversed the direction of their dancing. Now, they danced as if the sun were due west, in the direction of the station). In fact, of course, it was only somewhat west of south. They should have danced as if one had to fly with the sun on the left, which is how bees with full experience of the solar ephemeris at that locale in fact danced, but, because they relied on a (now incorrect) inference, based on limited experience, they danced as if one had to fly straight into the sun.

In summary, the mechanism by which bees learn the solar ephemeris has built into it what is universally and enduringly true about what has to be learned. It seems likely that this knowledge is built in by means of a genetically specified function relating the sun's direction to the time of day. The function takes as its input (argument) the time of day signal from the circadian clock, and it gives as its output a cyclically varying signal that specifies the compass direction of the sun at that time of day. The function is such that no matter what the values of its parameters, the sun is predicted to be somewhere in the east in the morning and somewhere in the west in the afternoon. Just as the entrainment mechanism extracts from experience the information necessary to set a parameter of the circadian clock (its phase), a learning mechanism extracts from experience the information necessary to set the parameters of the ephemeris function so that it correctly predicts the local direction of the sun at any time of day. Again, this is

also not a general purpose learning mechanism. It is structured to extract a particular kind of information from a particular kind of input.

Learning a Language

The just rendered account of how bees learn the solar ephemeris is strikingly similar to a widely (though by no means universally) accepted theory of how humans learn a language (Baker, 2001; Chomsky, 1975, 1981; Lightfoot, 1999). According to this theory, humans are the only species that can learn a human language because they have a unique genetic endowment, which specifies a universal grammar. The universal grammar specifies the things that are universally true about human languages. The universal grammar, like the universal solar ephemeris function, has parameters. Different values for these parameters yield different languages. Learning the local language is a matter of extracting from what one hears people say the values of the parameters.

Because the vast majority of sentences we both hear and produce are unique (heard or produced only once in a lifetime), we are constantly in the position of the bees in the just recounted experiment. Our previous experience does not—in the absence of extrapolation—give sufficient information to specify an interpretation for what we hear and say, that is, a mapping between the prelinguistic representation of what was to be communicated and the utterance itself. It is the universal grammar together with the parameter values derived from local experience that provides the mapping. In the absence of the universal grammar, most utterances would not mean anything to most hearers, because none of the utterances they had heard in the past would be the same as the utterance just heard. Similarly, in the absence of a universal grammar, speakers would not be able to convert their prelinguistic thoughts into the unique utterances by which they generally communicate their thoughts, because they would never before have experienced the requisite meaning-utterance pairing. This is closely analogous to the case of the bee, which, in the absence of a universal solar ephemeris function, would have no way to estimate from its experience of the sun's direction at other times of day what its direction might be at the current time of day.

In summary, there is reason to think that both the learning of the solar ephemeris in bees and the learning of language in humans is a matter of extracting from experience the information necessary to specify the parameters of built in functions. Although the two mechanisms are similar at this level of abstraction, they are otherwise dissimilar. No one would suppose that the mechanism that enables a bee to learn the solar ephemeris would enable it to learn a human language. And only the most extreme proponent of the Whorfian view that all human knowledge depends on language would argue that the mechanism that enables us to learn a language enables us to learn where to look for the sun when we emerge from our house in the morning.

Conditioning

Conditioning refers to the learning that occurs in experimental paradigms designed to study the laws of association formation. In these paradigms, the experimenter programs the animal's experience in such a way that an outcome that it either wants or wants to

avoid (typically, food, in the first case, and shocks to the feet in the second case) is contingent on either a neutral stimulus (typically, the sounding of a noise or tone or the appearance of a luminous pattern) or on a response of the animal. When the outcome, which is called the reinforcer or unconditioned stimulus (US for short), is contingent only on the occurrence of a motivationally neutral stimulus (called the conditioned stimulus or CS for short), then the process is called Pavlovian conditioning. When the outcome is contingent on something the animal does, it is called instrumental or operant conditioning.

In associative approaches to learning, conditioning and learning are synonymous, because conditioning paradigms are those by which psychologists determine the laws of learning (Miller & Escobar, 2002). In the light of the examples just discussed, one might wonder whether it makes sense to speak of the laws of learning. If problem-specific computational principles are needed to understand everything from ants learning their location, bees learning the solar ephemeris, and humans learning language, should we expect that there exists a form of learning that does not make use of problem specific computational principles? Clearly, if it does exist, the place to look for it is in the results from conditioning experiments. Moreover, most contemporary work on the neurobiology of learning is predicated on the assumption that the learning that occurs in conditioning experiments is the foundation of most if not all learning (e.g. Fanselow, 1993; Hawkins & Kandel, 1984; Mayford, Abel, & Kandel, 1995; Usherwood, 1993).

Because associative theories of learning do not specify how an animal extracts information from its experience, descriptions of conditioned behavior commonly do not discuss the timing of conditioned responses. In those paradigms in which the time of reinforcement occurs at a fixed interval following an earlier event, like the onset of the CS, or the delivery of the preceding reinforcement, conditioned responses are timed so as to approximately coincide with the predicted reinforcement (Figure 2). The timing of their response is evidence that subjects in conditioning experiments extract from the flow of events the durations of the intervals between events. Like the direction of the ant's homeward track, this response timing is a direct manifestation in subsequent behavior of information that has been extracted from previous experience and carried forward in memory. It suggests that the learning that occurs in conditioning experiments might be profitably subject to the kind of problem-specific computational analysis that was applied to the preceding examples. Conditioned behavior might also be the result of problem-specific computations that extract relevant information from the subject's experience.

Although conditioning paradigms were not developed within a conceptual framework that emphasizes the problem-specific nature of the learning principles, they do in fact pose a specific kind of computational problem; they present the subjects with problems in multivariate, non-stationary time series analysis. In a conditioning paradigm, the next occurrence of the reinforcer is probabilistically dependent on the earlier occurrence of another event or events (a CS, a response, or both), so, from a computational perspective, the problem the subject must solve belongs to the domain of time series analysis. The problems posed are generally multivariate because the subjects must generally compute which of several experienced CSs predicts the US (Kamin, 1969; Rescorla, 1968; Wagner, Logan, Haberlandt, & Price, 1968). Finally, they require what is

called a non-stationary time series analysis, because the predictive relation between the CS and the reinforcer or between the response and the reinforcer often changes in the course of the experiment (for example, in experiments that have a conditioning phase, during which the CS is paired with the US, followed by an extinction phase, when it is not).

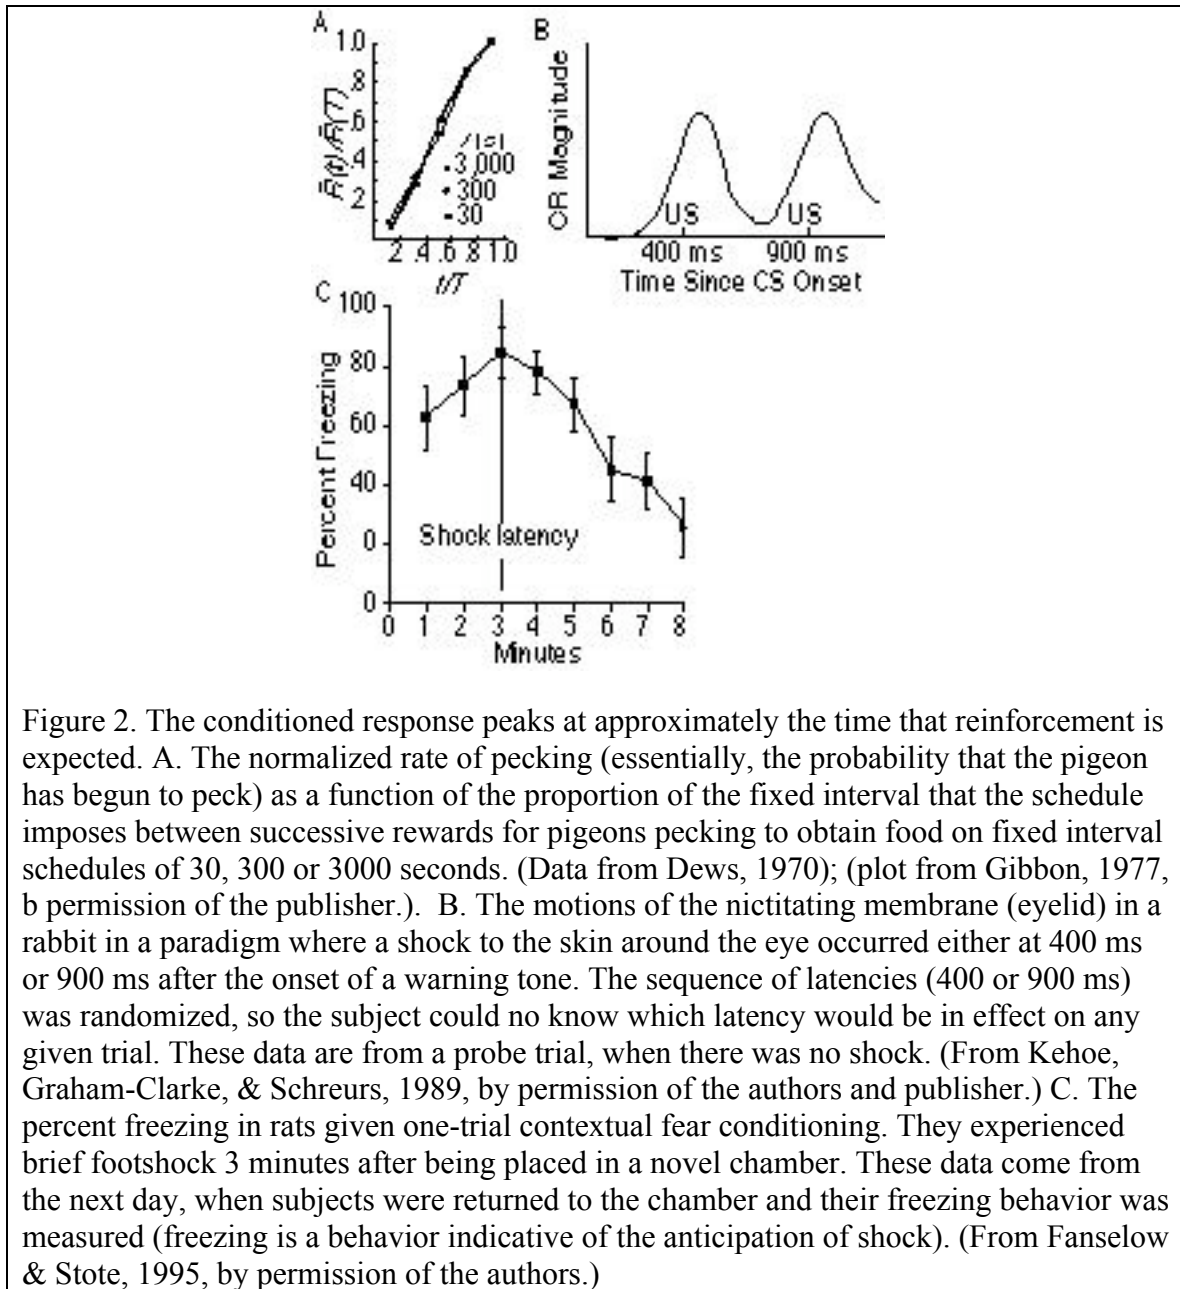


Figure 2. The conditioned response peaks at approximately the time that reinforcement is expected. A. The normalized rate of pecking (essentially, the probability that the pigeon has begun to peck) as a function of the proportion of the fixed interval that the schedule imposes between successive rewards for pigeons pecking to obtain food on fixed interval schedules of 30, 300 or 3000 seconds. (Data from Dews, 1970); (plot from Gibbon, 1977, by permission of the publisher.). B. The motions of the nictitating membrane (eyelid) in a rabbit in a paradigm where a shock to the skin around the eye occurred either at 400 ms or 900 ms after the onset of a warning tone. The sequence of latencies (400 or 900 ms) was randomized, so the subject could not know which latency would be in effect on any given trial. These data are from a probe trial, when there was no shock. (From Kehoe, Graham-Clarke, & Schreurs, 1989, by permission of the authors and publisher.) C. The percent freezing in rats given one-trial contextual fear conditioning. They experienced brief footshock 3 minutes after being placed in a novel chamber. These data come from the next day, when subjects were returned to the chamber and their freezing behavior was measured (freezing is a behavior indicative of the anticipation of shock). (From Fanselow & Stote, 1995, by permission of the authors.)

These considerations suggest that conditioned behavior might be profitably analyzed from the same computationally oriented, information-processing perspective that seems relevant to understanding the already discussed examples of problem-specific learning mechanisms. Gallistel and Gibbon (2000) have reviewed the conditioning

literature from this perspective, showing that this kind of analysis resolves several seemingly intractable paradoxes.

In their analysis, the multivariate problem is solved by means of a matrix computation that rests on the principle that independent random rates are additive. A random rate process is the simplest kind of process that generates recurring events. A random rate (or Poisson) process is the default assumption in time series analysis in the same way that a normal distribution is the default assumption in many other kinds of statistical analyses. If you know nothing about a US (reinforcing event) other than that n events have been observed in t amount of time, then you assume that it is a random rate process, because a random rate process is entirely characterized by a rate parameter, the number of events per unit of time. A random rate process is simpler from a statistical point of view than a process that generates a normal distribution, because to specify the latter you need two parameters, a mean and a standard deviation. A random rate process is also the simplest from an information processing point of view, because it has the property that the event is equally likely at any moment in time. The rate parameter determines what that likelihood is, but the likelihood is everywhere the same. Moreover, the process has no memory for its past: the expected time to the next occurrence of the event is the same no matter how long it has been since the last event. Knowledge of a rate parameter in a random rate process is the minimum knowledge it is possible to have about the time at which the next event will occur. The only thing less than this is no knowledge at all.

The above statements are true of random rate processes whether or not they are incorporated in the mechanism of conditioning, just as the laws of optics are true whether or not they are incorporated in biological lenses. However, in the model of the conditioning process proposed by Gallistel and Gibbon (2000), the mechanism that solves the multivariate aspect of the problem incorporates these principles just as faithfully as the vertebrate lens incorporates the laws of optics. Because random rate processes are additive, computing the true rates of US occurrence predicted by each of arbitrarily many assumed-to-be-independent CSs reduces to solving a system of simultaneous linear equations, in other words, to an exercise in matrix algebra. The formula, first given by Gallistel (1990), is:

$$\vec{\lambda}_t = \mathbf{T}^{-1} \vec{\lambda}_r,$$

where $\vec{\lambda}_t$ is the vector (the ordered list) of the (true) rates of US occurrence predicted by each CS acting in isolation, $\vec{\lambda}_r$ is the raw or uncorrected rate vector,

$$\vec{\lambda}_r = \left\langle \begin{array}{c} \frac{N_1}{T_1} \\ \frac{N_2}{T_2} \\ \vdots \\ \frac{N_m}{T_m} \end{array} \right\rangle$$

and \mathbf{T}^{-1} is the inverse of the temporal coefficient matrix:

$$\mathbf{T} = \begin{vmatrix} 1 & \frac{T_{1,2}}{T_1} & \dots & \frac{T_{1,m}}{T_1} \\ \frac{T_{2,1}}{T_2} & 1 & \dots & \frac{T_{2,m}}{T_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{T_{m,1}}{T_m} & \frac{T_{m,2}}{T_m} & \dots & 1 \end{vmatrix},$$

with iN_i is the cumulative number of USs in the presence of the i th CS, T_i is the cumulative amount of time that the i th CS has been present, and $T_{i,j}$ is the cumulative amount of time that i th and j th CS have been jointly present. Gallistel and Gibbon (2000) propose that subjects extract from their experience by means of counters and timers the requisite running total counts and durations (the N_i 's, T_i 's, and $T_{i,j}$'s) and make this computation in order to determine the rates of US occurrence predicted by each CS, and thence, on the assumption of independent causation (rate additivity), the rates of US occurrence to be expected from any combination of the CSs.

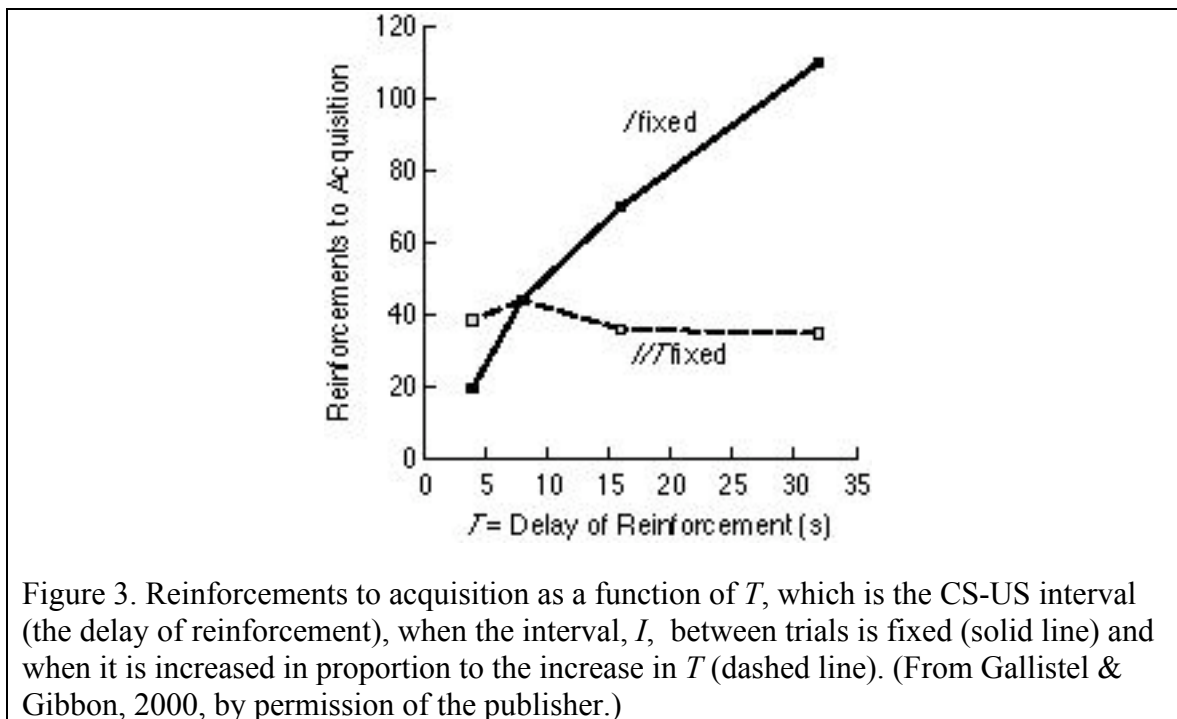
For present purposes, the only thing one needs to appreciate about this computation is that it is built directly on the objective (or analytic) properties of random rates. For that very reason, it is as problem specific as the computation with which we began, the integration of velocity with respect to time to obtain position with respect to time. No one would suggest that this matrix computation explains either how an ant learns where it is or how a human learns a language. This computation solves this problem and only this problem. It is but one example of the many problem-specific computations proposed by Gallistel and Gibbon (2000) to explain various experimentally determined properties of conditioning. If their analysis is any where near the mark, then the mechanisms that operate in conditioning are as problem specific as the ones already reviewed in other domains.

One may well ask, however, why one should take the Gallistel and Gibbon (2000) proposals seriously. Are there not explanations of conditioning results in terms of a general process associative theory, a theory (or theoretical framework) that incorporates no problem-specific principles? And are not these theories in some sense simpler and more parsimonious?

Gallistel and Gibbon's answer is that problem-specific models explain quantitative results that are deeply difficult for associative theories to explain, results that call into the question the basic assumptions of the associative framework. Many of these results come under the heading of the time scale invariance of conditioning. This empirical principle asserts that the time scale of a conditioning protocol does not matter; within broad limits, the results will be the same even if all the intervals in the protocol are changed by some multiplicative factor. Insofar as this principle holds, it implies that it is not the absolute durations of the intervals in the protocol that matters, it is the proportions among those durations. Scaling the protocol up or down does not change those proportions.

Take for example, the two proportions that exist in every classical conditioning protocol, the proportion between the ISI and the ITI. The ISI is the interstimulus interval, the interval between the onset of the CS and the onset of the US (the interval between the predictor event and the predicted event). The ITI is the intertrial interval, the interval between the offset of the US, which marks the end of one trial, and the next onset of the next CS, which marks the start of the next trial. The longer the intertrial interval is in relation to the interstimulus interval, the more precisely (relatively speaking) the CS predicts the occurrence of the US. By the principle of time-scale invariance, the results of a conditioning experiment should not change—other things being equal-- as long as the proportion between these two intervals does not change. In other words, what matters is the relative precision with which the CS predicts the US, how much it reduces the subject's uncertainty about when the US will occur (Gallistel, 2002, in press).

This principle appears to hold in the only conditioning paradigm in which it has been strongly tested (Figure 3). Increasing the ISI has no effect on the number of reinforcements that must be given before the conditioned response appears, provided the ITI is increased in proportion. So long as the relative precision with which the CS predicts the US is constant, the number of reinforcements that must be given to produce a conditioned response is constant.



The result in Figure 3 is difficult for associative theories to deal with because they assume that the temporal pairing of the CS and US is critical to conditioning. Temporal pairing seems like a simple notion, but what exactly does it mean? In associative theories it has meant that both the CS onset and the US onset must occur within a critical temporal window, or “trial.” If they both fall within that window (if, for example, they both occur on the same trial), then the two events are temporally paired; if they do not, then they are

not. If this is true, then conditioning cannot be time-scale invariant in the way just specified, because, as one scales up the temporal intervals in the protocol, there must come a point at which the ISI is longer than the critical window. Beyond that point, the conditioning process will no longer treat the CS and US as temporally paired, and conditioning will not occur. Put another way, the notion of temporal pairing in associative theorizing has always assumed that whether or not the associative process treats two events as paired depends on processes that have an intrinsic time scale of their own—temporal windows, trials, trace decay rates, state decay times, and so on (see, for example, Gluck & Thompson, 1987; Hawkins & Kandel, 1984; Rescorla & Wagner, 1972; Sutton & Barto, 1990; Tang et al., 1999; Usherwood, 1993; Wagner, 1981). The associative process is sensitive to the absolute durations of the intervals in a conditioning protocol, whereas the principle of time-scale invariance implies that it is only relative duration that matters.

A second example of the principle of time-scale invariance is seen in the lack of effect of partial reinforcement on the rate at which a conditioned response is acquired (Figure 4). It has been known since Pavlov (1928) that it is not necessary to deliver reinforcement every time the CS is presented. Even if one delivers the reinforcement on a randomly chosen fraction of the occasions on which it is presented, subjects nonetheless develop a conditioned response to the CS. In at least one commonly used conditioning paradigm, pigeon autoshaping, on average only one out of ten presentations of the CS may be reinforced, and still the subjects acquire. In fact, the number of reinforcements required to induce them to respond to the CS is no larger under these circumstances than it is when every CS presentation is reinforced (Figure 4). This result appears to directly contradict a basic assumption in associative theories, which is that reinforcement and non-reinforcement have opposing effects. Reinforcing the CS when it is presented strengthens the association; failing to reinforce it, either weakens the association, or, strengthens competing inhibitory associations. In either case, the behavioral effects of non-reinforcement oppose the behavioral effects of reinforcement. Therefore, intermixing an average of 9 non-reinforced presentations of the CS for every reinforced presentation, should greatly increase the number of reinforcements required to induce conditioned responding, which is contrary to experimental fact (Figure 4).

That partial reinforcement should not affect the outcome of a conditioning experiment (the number of reinforcements required to acquisition) is another manifestation of time-scale invariance. Interpolating non-reinforced presentations of the CS increases the average amount of exposure to the CS per reinforcement. Put another way, it reduces the rate of reinforcement in the presence of the CS. The more non-reinforced trials there are on average, the greater the reduction in this rate; hence the longer the subject must expect to wait for a reinforcement when the CS is present. In the previous example, however, we saw that how long a subject has to wait for its reinforcement in the presence of the CS does not matter in and of itself. What mattered was how long that wait was relative to the expected wait in the absence of the CS (Figure 3). The subject's estimate of the latter wait depends on how much unreinforced experience it has with the experimental chamber during the periods when the CS is not on, that is, during the intertrial intervals. The more unreinforced intertrial intervals it has experienced the lower its estimate of the background rate of reinforcement, the rate in the

absence of the CS. When we interpolate unreinforced presentations of the CS, we also interpolate an equal number of unreinforced intertrial intervals. Thus, if we ask, after r number of reinforced presentations of the CS, what is the ratio (proportion) between the expected wait when the CS is present and the expected wait when it is absent, the answer is unaffected by whether or not we have interpolated unreinforced presentations of the CS. The interpolation increases both waits by the same factor. Thus, the interpolation of unreinforced presentations should have no effect, as in fact it does not.




Figure 4. A. Time line showing a simple conditioning protocol. The elevations in the time line indicate the intervals during which a key on the wall of the pigeon's chamber was illuminated. The dots on the line at the end of CS presentations signify the delivery of food, which always came at the end of key illumination (whether or not the pigeon pecked the key). This Pavlovian procedure induces (classically conditions) key pecking. Reinforcement may be delivered every time the CS is presented (top line) or on some random fraction of the presentations (second line). B. Reinforcements to acquisition (dashed lines) and trials to acquisition (solid lines) as a function of the fraction of the CS presentations (trials) that terminate in reinforcement (food delivery). Each pair of solid and dashed lines is joined at the left, where the number of reinforcements equals the number of trials (a 1/1 schedule of reinforcement). The between curve variable is the

ratio of the interval between trials to the trial duration; the greater this ratio, the smaller the proportion of total time the conditioned stimulus (the key light) is on.

Whereas the time-scale invariance of the conditioning process directly challenges fundamental assumptions in associative theories of conditioning, it is a direct consequence of fundamental assumptions in a problem-specific approach. Problem-specific information-processing models assume that the learning mechanisms that extract relevant information from a given domain of experience are tuned to the deep principles governing that domain. From a problem-specific perspective, what the subject tries to extract from the temporal flow of events it experiences in a conditioning experiment is information that reduces its uncertainty about when the next event it cares about will happen (i.e., the next reinforcer). This is the kind of information that a time-series analysis tries to extract, which is why the mathematics of time-series analysis are assumed to be relevant to characterizing the underlying learning mechanism in conditioning. Time series analyses, like all statistical analyses, are scale invariant; they are unaffected by rescaling operations that increase or decrease all the inputs by some multiplicative factor. The following consideration shows that this must be true, without involving any actual mathematics: Computer programs that do statistical analyses, do not prompt the user for the units that go with the numbers in the data files. This implies that the units are irrelevant to the analysis, which implies that rescaling the numbers will have no effect, because changing the units from, for example, minutes to seconds, simply scales up all the numbers by a factor of 60. Thus, without going into mathematical details, it is possible to appreciate that a problem-specific, information-processing approach to understanding mechanisms of conditioning will be time-scale invariant. Hence, absent some overriding assumptions, it will predict the manifestations of time scale invariance in Figures 3 and 4.

Summary

The biological principle of adaptive specialization applies to learning and memory mechanisms just as much as it applies to other biological mechanisms. At the memory level, this suggests that we should look for the properties in the neurobiological mechanism of memory that we see in other good memory mechanisms. These properties are thermodynamic stability and high density.

We might also expect to find that the basic mechanism of memory, the mechanism that carries information forward in time, is domain independent, that is, it does not differ depending on the kind of information that is to be carried (visual information, auditory information, spatial information, temporal information, declarative information, episode information, biographical information, procedural information, etc.) Computer memories and DNA both have the function of carrying information forward in time, they both display the properties of thermodynamic stability and high density, and they both are domain independent. The action potential mechanism has the function of conveying information rapidly over long distances within the nervous system. It, too, is domain independent.

Finally, mechanisms that convey information are necessarily coding mechanisms. There must be principles that specify what it is about the state of the mechanism that conveys information about states of the world. Thus, another thing we can look for in the neurobiological mechanism of memory is the coding principle: how does the physical change that constitutes the memory represent information about the world?

In contrast to mechanisms of memory, there are reasons to think that learning mechanisms should be and are domain specific, just as are sensory organs. Learning involves extracting from raw sensory experience information about behaviorally important properties of the environment. Extracting this kind of information requires computations that are specific to the kind of information to be extracted and the sense data from which it is to be extracted. A very simple illustration of this principle is the dead reckoning mechanism, which enables subjects to learn their location from their experiences while getting to that location. The crucial sense data in this case are those that indicate the animal's velocity (its direction and speed of travel). The crucial computation is the integration of these velocity signals with respect to time to obtain signals that specify position.

This mechanism for extracting information from experience differs strikingly from the mechanism that learns the local solar ephemeris, which is the compass direction of the sun as a function of the time of day. At the heart of this mechanism is a genetically specified universal ephemeris function, with settable parameters. Learning the local ephemeris involves extracting from observations of the solar bearing at different times of day the requisite values of the parameters. When the parameters are correctly adjusted, this mechanism takes the time-of-day signal from the built-in circadian clock and generates a signal specifying the compass direction of the sun at that time.

The mechanism for learning the local solar ephemeris appears similar in broad outlines to the language learning mechanism envisaged by many linguists. At the heart of this mechanism is a genetically specified universal grammar, likewise with settable parameters. Learning to speak and understand the local language involves extracting from samples of that language the requisite values of grammatical parameters. When those parameters are correctly set, the subject speaks and understands the local language.

If problem-specific computational mechanisms are involved in everything from an ant's learning where it is, to a bee's learning where the sun is, and a human learning a language, one might reasonably suppose that learning is always mediated by one or more problem-specific mechanisms, whose structure suits them to solve one and only one class of problems. This supposition, however, conflicts with the long-standing tradition that an associative learning process mediates most learning. If all learning is mediated by problem-specific learning mechanisms, what is going on in Pavlovian and instrumental conditioning paradigms, which were created to study the laws of associative learning? It turns out that these paradigms all challenge subjects with a computationally coherent class of problems: multivariate, non-stationary time series analysis. Moreover, learning mechanisms specifically designed for this kind of problem deal successfully with quantitative property of conditioning—time scale invariance—that challenges the fundamental assumptions of associative theory.

If even Pavlovian conditioning is mediated by a problem-specific learning mechanism, then it is reasonable to conclude that learning processes are like other biological processes in that their structure is adapted to their function. There is a multiplicity of learning processes, each suited to solving a different kind of information-extraction problem.

References

- Baker, M. C. (2001). The atoms of language. New York: Basic Books.
- Braitenberg, V., & Schüz, A. (1991). Anatomy of the cortex. New York: Springer-Verlag.
- Carroll, S. B., Grenier, J. K., & Weatherbee, S. D. (2001). From DNA to diversity. London: Blackwell.
- Chomsky, N. (1975). Reflections on language. New York: Pantheon.
- Chomsky, N. (1981). Lectures on Government and Binding. Dordrecht: Foris.
- 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.
- Domjan, M. (1998). The principles of learning and behavior. Pacific Grove: Brooks/Cole.
- du Bois-Reymond, E. (1848). Untersuchungen über thierische Elektrizität. Berlin: Reimer.
- 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.
- Dyer, F. C., & Dickinson, J. A. (1996). Sun-compass learning in insects: representation in a simple mind. Current Directions in Psychological Science, 5, 67-72.
- Fanselow, M. S. (1993). Associations and memories: The role of NMDA receptors and long-term potentiation. 2(5), 152-156.
- Fanselow, M. S., & Stote, D. (1995). Temporal vs. associative accounts of the immediate shock deficit in contextual fear. Paper presented at the Annual Meeting of the Psychonomics Society. Nov. 11-12, Los Angeles.
- Frank, M. P. (2000). Lec. 3 & 4 in CIS 4930, [Web Page] [2002, 01/01].
- Frank, M. P. (in press 2002). Physical limits of computing. IEEE/AIP Computing in science and engineering.
- Gallistel, C. R. (1990). The organization of learning. Cambridge, MA: Bradford Books/MIT Press.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate and conditioning. Psychological Review, 107, 289-344.
- Gallistel, C. R. (2002, in press). Conditioning from an information processing perspective. Behavioural Processes.
- Gibbon, J. (1977). Scalar expectancy theory and Weber's Law in animal timing. Psychological Review, 84, 279-335.
- Gluck, M. A., & Thompson, R. F. (1987). Modeling the neural substrates of associative learning and memory: a computational approach. Psychol. Rev., 94(2), 176-191.
- Harkness, R. D., & Maroudas, N. G. (1985). Central place foraging by an ant (*Cataglyphis bicolor* Fab.): a model of searching. Anim. Behav., 33, 916-928.

- Hawkins, R. D., & Kandel, E. R. (1984). Is there a cell-biological alphabet for simple forms of learning? Psych. Rev., *91*, 375-391.
- Hebb, D. O. (1949). The organization of behavior: A neuropsychological theory. New York: Wiley.
- Hodgkin, A. L., & Huxley, A. F. (1952a). The components of membrane conductance in the giant axon of *Loligo*. Journal of Physiology (London), *116*, 473-496.
- Hodgkin, A. L., & Huxley, A. F. (1952b). Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. Journal of Physiology (London), *116*, 449-472.
- Hodgkin, A. L., & Huxley, A. F. (1952). The dual effect of membrane potential on sodium conductance in the giant axon of *Loligo*. J. Physiol., *116*, 497-506.
- Hodgkin, A. L., & Huxley, A. F. (1952c). A quantitative description of membrane current and its application to conduction and excitation in nerve. Journal of Physiology (London), *117*, 500-544.
- Hodgkin, A. L., Huxley, A. F., & Katz, B. (1952). Measurement of current-voltage relations in the membrane of the giant axon of *Loligo*. Journal of Physiology (London), *116*, 424-448.
- Hull, C. L. (1930). Knowledge and purpose as habit mechanisms. Psychol. Rev., *37*, 511-525.
- Judson, H. (1980). The eighth day of creation. New York: Simon & Schuster.
- 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.
- Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (2000). Principles of neural science. New York: McGraw Hill.
- Kehoe, E. J., Graham-Clarke, P., & Schreurs, B. G. (1989). Temporal patterns of the rabbit's nictitating membrane response to compound and component stimuli under mixed CS-US intervals. Behavioral Neuroscience, *103*, 283-295.
- Lightfoot, D. (1999). The development of language: acquisition, change and evolution. London: Blackwell.
- Mayford, M., Abel, T., & Kandel, E. R. (1995). Transgenic approaches to cognition. Current Opinion in Neurobiology, *5*, 141-148.
- Miller, R. R., & Escobar, M. (2002). Laws and models of basic conditioning. In C. R. Gallistel (Ed.), Stevens handbook of experimental psychology. 3rd ed. Vol. 3. Learning and motivation (Vol. 3, pp. 47-102). New York: Wiley.
- Pavlov, I. V. (1928). Lectures on conditioned reflexes: The higher nervous activity of animals (H. Gantt, Trans.). London: Lawrence & Wishart.
- 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., & 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.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., & Bialek, W. (1997). Spikes: Exploring the neural code. Cambridge, MA: MIT Press.

- Rolls, E. T., & Treves, A. (1998). Neural Networks and Brain Function. Oxford: Oxford University Press.
- Skinner, B. F. (1950). Are theories of learning necessary? Psychol. Rev., *57*, 193-216.
- Smolensky, P. (1986). Information processing in dynamical systems: foundations of harmony theory. In D. E. Rumelhart & J. L. McClelland (Eds.), Parallel distributed processing: foundations (Vol. 1, pp. 194-281). Cambridge, MA: MIT Press.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), Learning and computational neuroscience: Foundations of adaptive networks (pp. 497-537). Cambridge, MA: Bradford/MIT Press.
- Tang, Y.-P., Shimizu, E., Dube, G. R., Rampon, C., Kerchner, G. A., Zhuo, M., Liu, G., & Tsien, J. Z. (1999). Genetic enhancement of learning and memory in mice. Nature, *401*, 63-69.
- Turek, F. W. (1998). Circadian rhythms. Hormone Research, *49*, 109-113.
- Usherwood, P. N. R. (1993). Memories are made of this. Trends in Neurosciences, *16*(11), 427-429.
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), Information processing in animals: memory mechanisms (pp. 5-47). Hillsdale, NJ: Lawrence Erlbaum.
- 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., & Srinivasan, M. V. (1981). Searching behavior of desert ants, genus *Cataglyphis* (*Formicidae*, Hymenoptera). Journal of Comparative Physiology, *142*, 315-338.
- Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for search. IEEE Transactions on Evolutionary Computation, *1*(1).