

The Neuroscience of Learning: Beyond the Hebbian Synapse

C.R. Gallistel¹ and Louis D. Matzel²

¹Rutgers Center for Cognitive Science and ²Department of Psychology, Rutgers University, Piscataway, New Jersey 08854-8020; email: galliste@rucss.rutgers.edu

Annu. Rev. Psychol. 2013. 64:169–200

First published online as a Review in Advance on July 12, 2012

The *Annual Review of Psychology* is online at psych.annualreviews.org

This article's doi:
10.1146/annurev-psych-113011-143807

Copyright © 2013 by Annual Reviews.
All rights reserved

Keywords

LTP, spatial learning, Bayesian inference, information theory, cognitive map, geometric module

Abstract

From the traditional perspective of associative learning theory, the hypothesis linking modifications of synaptic transmission to learning and memory is plausible. It is less so from an information-processing perspective, in which learning is mediated by computations that make implicit commitments to physical and mathematical principles governing the domains where domain-specific cognitive mechanisms operate. We compare the properties of associative learning and memory to the properties of long-term potentiation, concluding that the properties of the latter do not explain the fundamental properties of the former. We briefly review the neuroscience of reinforcement learning, emphasizing the representational implications of the neuroscientific findings. We then review more extensively findings that confirm the existence of complex computations in three information-processing domains: probabilistic inference, the representation of uncertainty, and the representation of space. We argue for a change in the conceptual framework within which neuroscientists approach the study of learning mechanisms in the brain.

Contents

INTRODUCTION	170
ASSOCIATIVE LEARNING AND	
SYNAPTIC PLASTICITY	170
Effects of Interstimulus Interval and	
Intertrial Interval	171
Persistence	172
Reacquisition	172
Coding	172
REINFORCEMENT LEARNING ...	173
INFORMATION-PROCESSING	
DOMAINS	175
Probabilistic Inference	176
Representation of Uncertainty	177
Spatial Learning	178
CONCLUSIONS	192

INTRODUCTION

The theoretical frameworks with which we describe learning and memory have traditionally taken one of two forms. In the associative conceptual framework, the mechanism of learning cannot be separated from the mechanism of memory expression. At the psychological level of analysis, learning is the formation of associations, and memory is the translation of that association into a behavioral change. At the neuroscientific level of analysis, learning is the rewiring of a plastic nervous system by experience, and memory resides in the changed wiring.

When approached from the second perspective, the information-processing perspective, learning and memory are distinct mechanisms with different functions: Learning mechanisms extract potentially useful information from experience, while memory carries the acquired information forward in time in a computationally accessible form that is acted upon by the animal at the time of retrieval (Gallistel & King 2009). We review portions of the recent behavioral neuroscience literature, briefly from the first perspective, and more extensively from the latter perspective, focusing on neurobio-

logical systems that extract different kinds of information from different kinds of experience.

The distinction between the associative and information-processing frameworks is of critical importance: By the first view, what is learned is a mapping from inputs to outputs. Thus, the learned behavior (of the animal or the network, as the case may be) is always recapitulative of the input-output conditions during learning: An input that is part of the training input, or similar to it, evokes the trained output, or an output similar to it. By the second view, what is learned is a representation of important aspects of the experienced world. This representation supports input-output mappings that are in no way recapitulations of the mappings (if any) that occurred during the learning.

Before focusing on domain-specific learning, and most extensively on spatial learning, we briefly review some of the vast neuroscience literature on the two commonly proposed general-purpose associative learning mechanisms, Pavlovian conditioning (a.k.a. classical conditioning) and reinforcement learning (a.k.a. instrumental conditioning, a.k.a. operant conditioning). Historically, these theories are non- (or anti-) representational. They propose that the brain adapts behavior (its input-output mappings) to environmental circumstances without representing those aspects of the environment that make the behavior adaptive. Most contemporary cognitive science is, by contrast, representational; it assumes that brains construct a behaviorally useful representation of the experienced world through extensive computation.

ASSOCIATIVE LEARNING AND SYNAPTIC PLASTICITY

The hypothesis that the modification of synaptic transmission by experience mediates associative learning dates back to the elaboration of the concept of the synapse itself (Cajal 1894, Tanzi 1893). Hebb's (1949) influential statement of the hypothesis was that if a presynaptic neuron repeatedly played a role in firing a postsynaptic neuron, there ensued an enduring modification

of synaptic structure, such that activity in the presynaptic neuron became more likely to excite activity in the postsynaptic neuron. A snappier statement of this idea is that neurons that fire together wire together. Synapses that exhibit these properties are commonly called Hebbian synapses. Martin and colleagues (2000, 2002) review the arguments in favor of this hypothesis, which is widely accepted by psychologists, cognitive scientists, and neuroscientists.

The neurobiological process or phenomenon now most often identified with the Hebbian synapse is long-term potentiation (LTP). Recently, interest has focused on a form of LTP called spike timing-dependent plasticity (STDP; for a recent review, see Caporale & Dan 2008). In a variety of neural circuits, an enduring modification of synaptic transmission is produced by varying the timing of weak and strong synaptic inputs over a range of a few tens of milliseconds. The sign of the modification depends critically on the relative strength of stimulation and the timing of the two inputs. For some parameter values, transmission increases; that is, a presynaptic spike now produces a “potentiated” (i.e., larger amplitude or shorter latency) postsynaptic response. For other combinations, transmission decreases; that is, a presynaptic spike now produces a reduced postsynaptic response.

Most of the neurobiological literature on LTP focuses on its cellular and molecular mechanism. The relevance of this research to the neuroscience of learning depends on the hypothesis that links LTP to associative learning and to memory. The evidence for this link would be strong if the properties of LTP aligned closely with those of the associative learning process as revealed by behavioral experimentation. Here we review those properties and conclude that the alignment is poor.

Effects of Interstimulus Interval and Intertrial Interval

Behaviorally measured association formation depends on time parameters in a fundamen-

tally different way than does LTP. In LTP, differences of a few milliseconds to at most a few tens of milliseconds in the timing of the pre- and postsynaptic inputs are critical. This dependence is often cited in support of the linkage hypothesis (Quinn 2005, Thompson & Mattison 2009, Usherwood 1993). There is, however, nothing in the associative learning literature showing a dependence of association formation on event-timing differences measured in tens of milliseconds. The interstimulus intervals in behavioral experiments are orders of magnitude longer (seconds, minutes, and hours rather than milliseconds).

More fundamentally, there is no independent effect of the interstimulus interval [ISI, also known as the conditioned stimulus-unconditioned stimulus (CS-US) interval] in behavioral association formation: The number of trials to acquisition of a conditioned response in Pavlovian conditioning depends on the ratio of the CS-US interval to the US-US interval. The shorter the CS-US interval is relative to the US-US interval, the fewer the trials to acquisition (Gallistel & Gibbon 2000, Gibbon & Balsam 1981, Gottlieb 2008, Lattal 1999, Ward et al. 2012). The critical role of the CS-US/US-US ratio is dramatically shown by holding the CS-US interval constant and progressively shortening the US-US interval; there comes a point at which the association that forms is inhibitory rather than excitatory (Kaplan 1984). In short, there is no critical interstimulus interval for the behavioral phenomenon. Moreover, the more widely separated the instances of pairing, the more rapidly the association develops. The opposites are true for LTP: There is a critical interstimulus interval, which is orders of magnitude smaller than any interval relevant at the behavioral level. And, the more widely separated the instances of pairing, the weaker their cumulative effect (de Jonge & Racine 1985).

In defense of the linkage hypothesis, it may be argued that “This [disconnect] is only paradoxical, however, if [it is assumed that] CS-US associations occur online at the level of individual synapses. It is less problematic if the time-scale for information representation

in a given brain region is different from that pertaining to events as they happen” (Martin & Morris 2002, p. 610). This argument stipulates that the temporal properties of LTP do not explain the temporal properties of behaviorally measured association formation.

Persistence

Behaviorally measured associations can last indefinitely, whereas LTP always decays and usually does so rapidly. Its rate of decay is measured in hours or days (for review, see Abraham 2003). Even with extended “training,” a decay to baseline levels is observed within days to a week (e.g., Castro et al. 1989). An experiment by Power et al. (1997) highlights the lack of correspondence: They recorded changes in CA1-evoked responses in brain slices obtained from animals trained on a trace eyeblink conditioning task, which is dependent on the hippocampus for its behavioral expression. Potentiated postsynaptic responses were observed at 1 hour, but not at 24 hours after training. By contrast, the learned eye-blink response remains intact for weeks or months. Again, it may be argued that “it would be premature to reject synaptic plasticity as a memory mechanism merely for this reason [lack of sufficient persistence]. Hippocampal LTP may need only last long enough (a few weeks perhaps) to permit completion of a slower neocortical consolidation process” (Martin & Morris 2002, p. 610). This argument stipulates that the persistence of LTP does not explain the persistence of associative learning.

Reacquisition

Although behavioral evidence for the presence of an association can generally be obtained months and even years after its establishment, the strength of the conditioned response does commonly decline somewhat with time. And, of course, the learned response may be weakened by extinction and/or counter-conditioning. Both forgotten and extinguished conditioned responses exhibit facilitated reacquisition; that is, they are relearned more efficiently than

when they were initially acquired (e.g., Napier et al. 1992; for review, see Miller et al. 1986). Following its decay to baseline, LTP is neither more easily induced nor more persistent than it was after previous inductions (de Jonge & Racine 1985).

Coding

Perhaps most importantly, the hypothesis that a change in synaptic transmission is the mechanism of memory does not address the coding problem. The encoding of the temporal intervals in conditioning protocols routinely occurs (Arcediano et al. 2003, Barnet et al. 1996, Blaisdell et al. 1998, Burger et al. 2001, Cole et al. 1995), probably before the emergence of the conditioned response (Balsam et al. 2006, 2010; Balsam & Gallistel 2009). More tellingly, the sign (excitatory or inhibitory) and rate of association formation depend on the ratio between the expectations of two intervals in the protocol (the CS-US and the US-US intervals; see Gallistel & Gibbon 2000, Ward et al. 2012), which suggests that the encoding of temporal intervals may be a precondition for the appearance of conditioned responses. Thus, the mechanism that mediates associative learning and memory must be able to encode the intervals between events in a computationally accessible form. There is no hypothesis as to how this could be accomplished through the modification of synaptic transmission.

The lack of suggestions in the literature about how Hebbian synapses might encode the durations of intervals reflects a more general failing of the associative conceptual framework when viewed from the perspective of cognitive science: It may explain reflex modification phenomena, but it does not explain the learning of behaviorally important facts and the formation of data structures. It fails to address the question of how facts about the experienced spatio-temporal environment may be carried forward for indefinite periods of time to inform subsequent behavior in ways not foreseeable when the facts were learned. The

neuroscientific literature on the representation of space, to which we soon turn, shows that the brain carries in memory the learned geometry of the experienced environment in a way that makes this acquired information accessible to computation in small fractions of a second.

The failure to address the coding problem would not count against the hypothesis that links LTP to memory if “what gets encoded and how is an emergent property of the network in which this plasticity is embedded, rather than of the mechanisms operating at the synapse in isolation” (Martin et al. 2000, p. 650). This appeal to emergent properties stipulates that the properties of LTP do not explain the essential property of a memory mechanism, the ability to store information in a computationally accessible form (Gallistel & King 2009). That most basic property is said to reside in “the network.” One naturally asks where it resides and how that storage is implemented. Does the claim that the storage of information is an emergent property imply that we are never to have answers to these questions?

In summary, if synaptic LTP is the mechanism of associative learning—and more generally, of memory—then it is disappointing that its properties explain neither the basic properties of associative learning nor the essential property of a memory mechanism. This dual failure contrasts instructively with the success of the hypothesis that DNA is the physical realization of the gene. This linkage hypothesis asserts that DNA is the molecule that stores hereditary information and makes it accessible to orchestrate ontogeny and much else. The structure of the molecule explains not only its ability to store information but also how copies of it may be made. There is no need to appeal to elusive (and possibly illusive) emergent properties in support of this linkage hypothesis. This explanatory power is a major reason why the hypothesis that links the gene to DNA is so much more compelling than the hypothesis that links LTP to associative learning and to memory.

In **Table 1**, we catalog the discrepancies between the properties of LTP and the properties of associative learning.

REINFORCEMENT LEARNING

Historically, two different association-forming processes have often been posited, one dependent only on temporal contiguity, the other on response-contingent reinforcement (and temporal contiguity). The latter process is often called instrumental conditioning, to distinguish it from the former, which is called Pavlovian or classical conditioning. In the traditional conception of the effects of reinforcement, there was no representation of the reinforcement. The reinforcement “stamped in” an association between a stimulus and the response that produced the reinforcement (Hull 1952), but neither the reinforcement nor its being a consequence of the response was represented. In the quite different contemporary formulation, which has been strongly influenced by theoretical work on reinforcement learning in computer science (Sutton & Barto 1998), reinforcement history is explicitly represented by a value variable associated with the response: A temporal-difference learning algorithm computes the value of an action in a given situation (Dayan & Daw 2008, Redish et al. 2007). The estimated value is updated after each performance of the response in proportion to an error term, which is the difference between the obtained reinforcement and the predicted reinforcement.

Neurobiological support for these models is found in the similarity between the value prediction error term and the signals observed in dopamine neurons following reinforcement and nonreinforcement. There are several recent, theoretically oriented reviews of the relevant literature (Berridge 2012, Dayan & Daw 2008, Flagel et al. 2011, Schultz 2006, Zhang et al. 2009). The general finding is that dopaminergic neurons in the basolateral diencephalon fire in response to events that occur at unpredictable times. If the time of reinforcement (the US) may be predicted by reference to an earlier temporal “landmark,” then dopaminergic neurons do not fire. Rather, they fire in response to the occurrence of the landmark (the CS), whose time of occurrence is

Table 1 Disparate properties of LTP and associative learning

Property	Hebbian LTP	Associative learning
Coding	Not implemented by LTP itself: an “emergent property” of circuits	Depends on encoding of temporal intervals, stimulus properties, and stimulus relationships
Necessary CS-US relation	Close temporal contiguity	Contingency
Form of learned output	Recapitulative: When stimulus recurs, output recurs	Anticipatory: Learned behavior usually differs from behavior during learning
Critical ISI	1–100 ms	None: Rate of conditioning is inversely proportional to ISI/ITI
Effect of ITI	The longer the ITI, the weaker the LTP	The longer the ITI relative to ISI, the faster and stronger the learning
Induction kinetics	Expression requires tens of seconds to minutes	Behavioral expression is immediate, <1 s after induction
Acquisition function	Requires repetition	Often complete within single trial
Persistence	Hours–weeks	Months–years (up to a lifetime)
Reacquisition	Not facilitated by previous acquisition	Facilitated by previous acquisition
Context learning	Not consistent with ISI requirement	Ubiquitous and fundamental
Long delay and trace conditioning	Seemingly incompatible with ISI requirement	Easily attained
Cue competition (blocking, overshadowing, etc.)	Not explained by properties of LTP	Ubiquitous and fundamental

Note: Some of these properties are not discussed in text; see Matzel & Shors (2001) and Gallistel & King (2009) for full discussion. For temporal pairing versus contingency, dependence of associative learning on ISI/ITI, and cue competition, see Gallistel & Gibbon (2000), Balsam & Gallistel (2009), Balsam et al. (2010), and Ward et al. (2012). Abbreviations: CS-US, conditioned stimulus-unconditioned stimulus; ISI, interstimulus interval; ITI, intertrial interval; LTP, long-term potentiation.

itself unpredictable. If the US fails to occur at the predicted time, the neuron fires. The neuron also fires if the US occurs at an unexpected time in relation to the CS. Thus, contemporary reinforcement learning theory assumes that the duration of the previously experienced CS-US interval resides in memory, where it forms the expectation against which a currently experienced CS-US interval is compared. The comparison between present experience of the CS-US interval and the information about past intervals stored in memory is on the causal pathway from a reinforcing event to the firing of dopaminergic neurons elicited by that event.

The enduring appeal of antirepresentational associative theory has been its neurobiological transparency: It is easy to imagine that the formation of an associative bond is physically realized by a change in synaptic transmission.

Conceptually, both are simple conductive connections. In associative learning theory, the associative bond does not represent an aspect of the experienced world, so our inability to specify how changes in synaptic transmission encode facts is not a problem. The convergence of behavioral and neuroscientific evidence on the conclusion that the coding of temporal facts (interval durations) is an essential feature of both Pavlovian and reinforcement learning suggests that the antirepresentational form of associative theorizing may need to be abandoned. If so, we must now face squarely the unanswered question as to the physical realization of the neural memory mechanism that stores simple abstract experiential facts, such as durations, distances, directions, and probabilities, in a structured form and makes them accessible to computation on a millisecond time scale.

INFORMATION-PROCESSING DOMAINS

In associative learning theory, the brain rewires itself so as to perform better in the experienced world, but it does not represent what it is about the world that makes the new performance better suited to it. When so conceived, the association-forming process may be mathematically modeled, but it is not the physical realization of a computation (or a memory). The information-processing framework, by contrast, is closely allied to the computational theory of mind, which holds that a necessary level of analysis in connecting neuroscience to behavioral phenomena is an analysis of the computations that the brain performs in extracting behaviorally useful information from raw experience (Marr 1982). On this theory, to understand the operations of the mind/brain, we must understand what aspects of the experienced world the brain represents (the representational question), how it represents them (the encoding question), how it computes that representation from the relevant aspects of its experience (the computational question), and how it translates its representations into behavior (the performance question).

Framing learning problems as computational problems leads to the postulation of domain-specific learning mechanisms (Chomsky 1975, Gallistel 1999) because no general-purpose computation could serve the demands of all types of learning. Some computations are broadly useful whereas others are only useful in a single context. However, they all apply the primitive operations of arithmetic and logic to different combinations of inputs to achieve different results.

Framing learning as the problem of how the brain computes a behaviorally useful representation of the experienced world more or less eliminates any distinction between perception and learning. In the study of perception, it is understood that an understanding of the physical and mathematical principles operative in a domain is a precondition for psychological and neuroscientific understanding of how

the brain functions in that domain. You cannot understand vision without understanding the rudiments of geometric optics. Similarly, the information-processing approach to learning mechanisms requires an understanding of the rudiments of the different domains in which different learning mechanisms operate. In the balance of this review, we consider three domains: probabilistic inference, the representation of uncertainty, and the representation of space. In each domain, we review the rudiments before focusing on neuroscientific findings relevant to the first two of Marr's questions: What is represented, and how is it represented?

When it is assumed that the neuroscience of learning is the neuroscience of synapse modification, then the study of processes that modify synaptic transmission is naturally conceived of as the study of the cellular and molecular mechanism of learning. But if learning is the result of domain-specific computations, then studying the mechanism of learning is indistinguishable from studying the neural mechanisms that implement computations. Although there is a large body of theoretical work in computational neuroscience, there is as yet no consensus about foundational questions, such as:

1. How is information encoded in spike trains?
2. What are the primitive computational operations in neural tissue?
3. Are they implemented at the network level, the molecular level (intracellularly), or both?
4. What cellular and/or molecular mechanisms implement the arithmetic operations?
5. What mechanism implements memory (the storage of information in a computationally accessible form; see Gallistel & King 2009)?
6. What mechanism implements variable binding in memory? (For an explanation of variable binding and its importance in computation, see Gallistel & King 2009.)
7. What mechanism implements data structures in memory?

Thus, at this stage of our science, neuroscientific findings bear strongly on representational questions in learning—on what is learned—but they do not yet give us a cellular and molecular understanding of underlying computational mechanisms.

Our review of the neuroscience of domain-specific learning mechanisms begins with the neural mechanism of Bayesian inference because it is an example of a broadly applicable complex computation. Its relevance to perception is now well understood, but it also applies to learning, because learning the state of the world is an inferential process. Gallistel (2012) models extinction as Bayesian change detection. This treatment of a basic issue in traditional learning theory explains quantitatively the partial reinforcement extinction effect,¹ which has resisted principled explanation for more than half a century. A second reason for beginning with the neuroscience of probabilistic inference is that there is interesting recent work on the neural mechanism of marginalization, which is an essential component of Bayesian inference.

Probabilistic Inference

Rudiments: Bayes rule. The Bayesian computation mediates probabilistic inference about the state, w , of some aspect of the world by taking the product of a likelihood function and a prior probability distribution:

$$L(w|\mathbf{D}, \pi(w)) = L(w|\mathbf{D})\pi(w).$$

The prior distribution, $\pi(w)$, represents the probability of the different possible states in the light of previous or extraneous evidence. The likelihood function, $L(w|\mathbf{D})$, represents

the likelihood of those states in the light of some new data (or event or signal), symbolized by \mathbf{D} , that carries information about that state of the world. The product of the two functions weights each prior probability by the corresponding likelihood, giving the relative likelihoods of the different possibilities “all considered.” When normalized so that it integrates to one, the product is called the posterior probability distribution.

Rudiments: Likelihood. From a neurobiological perspective, a likelihood function is a neuronal firing pattern viewed backwards, viewed, that is, from the brain’s perspective rather than from the perspective of the experimenter (cf. Rieke et al. 1997), which is why likelihood is sometimes called reverse probability. The experimenter, who knows the experimentally given stimulus (the relevant state of the world), determines by experiment the probabilities of the different possible neuronal responses (different numbers of spikes) and plots them as a neuronal tuning curve, a staple of experimental systems neuroscience. The tuning curves for an array of neurons of the same class, for example, the simple cells in the primary visual cortex (V1), determine the firing pattern across an array of neurons of that class. However, from the brain’s perspective, this firing pattern is the given; what it must infer are the relative likelihoods of the different plausible stimuli, i.e., what it was out there in the world that produced this pattern of firing.

Rudiments: Marginalization. To assess the risks attendant on different possible decisions, the brain needs to represent the entire likelihood function, not just the most likely value of w . The problem it almost always faces is that several different aspects of the world affect the firing pattern. For example, the firing of the so-called simple neurons in the primary visual cortex (V1) is jointly determined by (among several other factors) the location, orientation, and contrast of the image on the retina. Put more formally, tuning curves, hence also the likelihood functions, are almost always

¹Partial reinforcement during training increases the number of unreinforced trials or responses required to extinguish the learned response. This is paradoxical in associative learning theory because the unreinforced trials during training should weaken net excitatory associative strength at the end of training. Fewer nonreinforced trials should then be required to reduce this weak association to insignificance, whereas in fact the number of unreinforced trials required for extinction increases in proportion as the pre-extinction probability of reinforcement decreases.

multidimensional. For any one decision, the brain commonly needs a one-dimensional likelihood function, a function, for example, that gives the likelihoods of different possible orientations, regardless of contrast and location. To obtain a one-dimensional likelihood function, it must marginalize the multidimensional likelihood function; that is, it must “integrate out” the effects of the “nuisance” parameters. (In this example, they are the location and contrast in the light pattern.) Metaphorically, marginalization is a bulldozer that moves along one dimension of a multidimensional likelihood function, piling up the likelihood against an orthogonal wall.

Neural implementation. Beck and colleagues (2011) show that combining two widely observed properties of neuronal stimulus-response functions—divisive normalization and quadratic nonlinearities—gives a neurobiologically plausible implementation of marginalization. Divisive normalization is a form of lateral inhibition in which the response of one neuron in a class of neurons (e.g., the simple cells in V1) is suppressed in proportion to the inverse of the sum of the responses of the other neurons in its class (see, e.g., Heeger 1992, Olsen et al. 2010, Simoncelli & Heeger 1998). Quadratic nonlinearities occur when the inputs to a neuron combine multiplicatively rather than additively (see, e.g., Andersen et al. 1985, Galletti & Battaglini 1989, Groh et al. 2001, Werner-Reiss et al. 2003).

Beck et al. (2011) emphasize the broad range of applications of marginalization, from coordinate transformations to causal inference. As they stress, it is a key operation in Bayesian inference. According to much contemporary thinking in cognitive science, Bayesian inference is everywhere in cognition, from perception to learning and causal reasoning (Chater et al. 2006, Griffiths et al. 2010).

Representation of Uncertainty

Probabilistic inference plays a central role in the construction of useful representations of the

experienced world because there is a complex, noisy, and ambiguous relation between the behaviorally important properties of the world and the first-order neural signals from which the brain must infer the states of the world. The inferences to be drawn from sensory input are for that reason uncertain to varying degrees. The quantification of this uncertainty through information-theoretic computations complements Bayesian inference.

Rudiments: Uncertainty = entropy = available information. A counterintuitive aspect of information theory is that information and uncertainty are two words for the same quantity. The objective amount of uncertainty about some aspect of the world, that is, the range of possibilities and their probabilities, is the source information, also called the available information. It puts an upper limit on the amount of information that a neural signal (or correlated event, or variable, or memory) can convey about that aspect of the world. Intuitively, the more uncertainty there is about something, the more there is to learn, that is, the more information to be gained. If there is no uncertainty, then there is nothing to learn, that is, no information to be gained.

A probability distribution specifies the probabilities of a set of mutually exclusive and exhaustive possibilities. The possibilities are the support for the distribution. Signals, signs, and events carry information insofar as they may be used to narrow the range of plausible possibilities. A probability distribution is an example of structured information: It links the symbols for the possibilities to the symbols for their relative frequencies.

The entropy of a probability distribution measures the amount of uncertainty. If the probability distribution is in the brain’s representation of an empirical variable, then its entropy measures the brain’s uncertainty about the value of that variable. A signal or predictor event is informative about that value to the extent that it reduces the entropy of this distribution, because information is defined as the reduction of uncertainty (Shannon 1948).

The formula for computing the entropy of a distribution is a probability-weighted sum, as are the formulae for mean, variance, and the still higher moments:

$$\begin{aligned}
 H &= \sum_{i=1}^{i=n} p_i \log(1/p_i) \text{ the entropy of the distribution} \\
 \mu &= \sum_{i=1}^{i=n} p_i w_i \text{ the mean (1st moment)} \\
 \sigma^2 &= \sum_{i=1}^{i=n} p_i (w_i - \mu)^2 \text{ the variance (2nd moment about the mean)} \\
 \gamma &= \sum_{i=1}^{i=n} p_i (w_i - \mu)^3 \text{ the skew (3rd moment about the mean)}
 \end{aligned}$$

where i indexes the possibilities that constitute the support for the distribution. The log of $1/p_i$ measures the information attendant on the occurrence of w_i , the i^{th} possibility. The summation weights each such amount by the relative frequency of its occurrence, that is, by p_i .

Rudiments: Contingency. An important advance in our understanding of associative learning came from experiments demonstrating that the emergence of a conditioned response depends not on the temporal pairing of two events but rather on the contingency between them. Eliminating the contingency while preserving the temporal pairing prevents the emergence of a conditioned response (Rescorla 1967, 1968). Thus the simple contiguity of events is insufficient to support learning. This is another case in which the Hebbian properties of LTP fail to explain the properties of associative learning: LTP is driven by temporal pairing; association formation is driven by contingency.

Neuroscientific evidence. The importance of this insight to our understanding of the neurobiology of associative learning is shown by the finding that signals in the mesolimbic dopaminergic neurons encode the probability and uncertainty of reinforcement (Fiorillo et al. 2003) and by the more recent discovery

that the response of neurons in the amygdala to reward-predicting stimuli depends on the contingency between the stimulus and the reward rather than on their temporal pairing (Bermudez & Schultz 2010). Thus, to understand the neurobiology of associative learning, a measure of contingency is needed.

Information theory provides a generally applicable measure:

$$C_{YX} = I_{YX}/H(p(X)) \neq C_{XY} = I_{YX}/H(p(Y)),$$

where I_{YX} is the mutual information between variables Y and X , C_{YX} measures the extent to which Y is contingent on X , and $H(p(X))$ is the entropy of the distribution of X , which distribution is symbolized by $p(X)$. The mutual information is the sum of the entropies of the individual distributions minus the entropy of their joint distribution:

$$I_{YX} = H(p(X)) + H(p(Y)) - H(p(X \times Y)),$$

where $p(X \times Y)$ symbolizes the joint distribution. Intuitively, the information-theoretic measure of contingency quantifies the extent to which knowledge of a putative predictor (Y , the CS or a response) reduces the uncertainty about when reinforcement (X) will occur.

Given the evidence from both behavior and neuroscience that contingency is fundamental, an important challenge for further neuroscientific investigation is to discover the mechanisms that represent distributions, compute their entropies, and measure the contingencies between events.

Spatial Learning

For decades, psychologists, cognitive scientists, and neuroscientists with an empiricist bent resisted the assumption that the mind/brain explicitly represents anything (Brooks 1991, Chemero 2011, Edelman 1989, Elman & Zipser 1988, Hull 1930, Markman & Dietrich 2000, Rumelhart & McClelland 1986, Shastri & Ajjanagadde 1993, Skinner 1938, Smolensky 1986), let alone aspects of experience as far removed from sense data as probability, uncertainty, time, and space.

There is, however, a large behavioral literature implying that learned representations of spatial locations and directions underlie animal navigation, including the navigation of many insects (for reviews, see Cheng 2008, Cheng et al. 2007, Collett & Collett 2004, Collett & Graham 2004, Gallistel 1990, Legge et al. 2010, Menzel et al. 2005, Merkle & Wehner 2008, Sommer et al. 2008, Wystrach et al. 2011). Another substantial literature implies the representation of time-of-day, time-of-month, time-of-year, and temporal duration and direction (Antle & Silver 2005, Bouton & Garcia-Gutierrez 2006, Budzynski & Bingman 1999, Buhusi & Meck 2005, Crystal 2001, Denniston et al. 2004, Gwinner 1996, Matzel et al. 1988, Meck 2003, Savastano & Miller 1998, Zhang et al. 2006), but there is not space to review that literature here.

Fundamentals of navigation and spatial representation. To assume that animals represent space is to assume that the brain has one or more spatial coordinate systems that encode locations in one or more frames of reference. It also assumes brain mechanisms for estimating distance and direction. Without a mechanism that implements a system of coordinates, there is no way to represent location. Without mechanisms for estimating direction and distance, there is no way to assign to a notable point in the environment a vector representing its location. In short, the postulation of spatial representations assumes the existence of nontrivial, genetically specified, purpose-specific representational machinery. This machinery does the spatial learning when it constructs a representation of the geometry of the experienced environment and tracks the animal's position and heading within that representation.

Frames of reference. A behaviorally useful coordinate system for representing location and/or direction must be anchored to a frame of reference. Coordinates are symbols that represent locations. Typically, they are vectors, ordered pairs (or triplets) of numbers that are subject, as ordered pairs, to some mathematical

operations such as addition and subtraction. A frame of reference is established when at least two of these vectors are assigned a referent in a physically instantiated space: this vector refers to that place or that direction. Assigning referents for at least two vectors establishes referents for all possible vectors (all the possible location symbols within a given framework). Changing the frame of reference changes which symbols refer to which locations. Symbols carry information forward in time (Gallistel & King 2009). Their physical realization in neural tissue is as yet unknown (unless one imagines that altered synaptic conductances can somehow function as symbols). Signals carry information from place to place. In the brain, information is carried over long distances by spikes (nerve impulses). As we will see, changes in the frame of reference are common in the neurons whose firing signals spatial locations and directions.

Frames of spatial reference fall into two broad classes, egocentric (self centered) and allocentric (other centered), depending on whether the system of coordinates is anchored to a part of the animal's body or to an aspect of the environment. Prominent among the egocentric frameworks are the eye-centered and head-centered frameworks. Prominent among the allocentric frameworks are the geocentric (earth-centered), enclosure-centered, object-centered, and array-centered frameworks.

A well-established behavioral result, to be borne in mind when assessing the neurobiological results, is that animals of widely diverse species maintain a geocentric orientation: a sense of their orientation (and location) in a framework anchored to an indefinitely extended surrounding environment (Baird et al. 2004, Douglas 1966, Dudchenko & Davidson 2002, Etienne et al. 1986). Mammals are compass-oriented even when they have no immediate sensory basis for this orientation. Their compass sense is based on inertial dead reckoning, not on the earth's magnetic field; that is, it is based on integrating the angular velocity signals from the semicircular canals in the ear. This integration rests on an implicit commitment to the principle that direction

(angular position) is the integral of angular velocity.

An early and striking manifestation of rats' compass orientation came in experiments designed to determine the cues that a rat uses in navigating a familiar maze (Carr 1917). Rats were trained to run rapidly through a complex maze, inside a square enclosure of heavy black curtains, within a large laboratory room. Between trials, the rats were kept in home cages at the other end of the room, outside and some distance from the curtained enclosure. Running trials in complete darkness had little effect on performance, as did blinding the rats, deafening them, or rendering them anosmic. On the other hand, rotating the maze and the surrounding curtain enclosure by 90° produced a profound disruption of maze performance, even though the maze itself and its relation to the perceivable surroundings (the black curtains) were in no way altered.

One of the present authors observed a similar effect in a similar experiment, again with rats (Margules & Gallistel 1988). The rats were trained to find buried food at previously demonstrated locations in a rectangular box with high walls and prominent, multimodally distinctive landmarks in the corners. The landmarks were intended to distinguish one end of the box from the other. Between trials, the rats were kept in a cage elsewhere in the room. Rotating the experimental box between trials within a normally lighted room noticeably upset them, causing freezing and other signs of fear. Despite the high walls, which prevented their seeing anything but the ceiling of the room when in the box, they were aware of and greatly perturbed by the change in the geocentric orientation of the test box. On the other hand, under red light (complete darkness for the rat) and after slow rotation of their cage for a few minutes before they were transferred to the box, which destroyed their inertial orientation, they were no longer perturbed by rotation of the test box in the room, because they could no longer detect it.

Our intent in rotating the box within the room was to force the rats to rely on the

prominent landmarks in the box to distinguish one end of the box from the other. In this, we failed. As in earlier experiments, when geocentrically disoriented, the rats ignored the corner landmarks when digging at what they took to be the location of the buried food, with the result that half the time they dug at the rotationally equivalent location (Cheng 1986, Gallistel 1990, Margules & Gallistel 1988). A rotationally correct location is correct except for a 180° rotation of the box; it is the location one digs at when one is misoriented within the rectangle. In other words, when the geometry of the test box limited the possible reorientations to two, rats consistently failed to use prominent landmarks in the corners to establish a unique (and correct) orientation.

The evidence for the maintenance of geocentric orientation does not imply that animals do not rely on more local frames of reference when navigating within enclosed spaces. Neurobiological results on place and head direction cells show clearly that they do use these local coordinate frameworks. However, unless the animal is geocentrically disoriented before placement in the enclosure, the geocentric orientation of the enclosure itself is also represented, even when this may not be apparent in the firing of head direction cells. This representation is a basis for the subjective polarization of symmetrical enclosures, such as rectangles and cylinders. Rotational confusions in symmetric enclosures (Cheng 1986) are observed only when the rats are geocentrically disoriented by slow rotation in the dark. When subjects enter an enclosure with their geocentric orientation intact, their geocentric orientation polarizes the enclosure, establishing for the animal which way is which within that enclosure. When an intrinsically polarized enclosure—one without rotational symmetries—has been rotated, then the animal's geocentric orientation enables it to detect and respond to that rotation.

The midbrain's capacity to integrate the angular velocity signal from the vestibular system so as to maintain the geocentric orientation by inertial means explains rats' remarkable sensitivity to changes in the geocentric orientation

of experimental closures (Angelaki et al. 2010, Rochefort et al. 2011). This capacity enables rats to carry the directional parallel from their cage and the larger room into test enclosures that eliminate or greatly restrict sensory access to the larger space. To “carry a parallel” is to preserve a directional axis when going from one part of the world to another (or one part of a map to another). Doing so is essential to dead reckoning, which is an essential aspect of navigation, map construction, and landmark recognition.

Dead reckoning, also known as path integration, plays a fundamental role in animal navigation (Cheung et al. 2008, Collett & Graham 2004, Gallistel 1990, Loomis et al. 1998, Wehner & Srinivasan 2003, Wittlinger et al. 2006). It is the reckoning of a new position and heading from an old position by summing successive intervening displacements and changes in heading to obtain the net change in position and heading.

Dead reckoning provides the animal with a moment-to-moment representation of its location and orientation on its cognitive map. The mechanism that mediates dead reckoning is a learning mechanism, by means of which the animal learns where it is. Diverse species of animals rely strongly on this dynamic representation of their position and heading, as did marine navigators until the very recent advent of the global positioning system. When a rat has learned to run a complex maze rapidly and the experimenter then shortens the paths, the rat runs full tilt into the walls at the end (Carr 1917). If it is an elevated maze, the rat runs off the end of the shortened segments into thin air (Dennis 1932). These results are analogous to the many shipwrecks caused by faulty dead reckoning. The rat, like the mariner, only looks (or feels) for landmarks when its dead reckoned position on its map approaches a mapped boundary or waypoint.

Dead reckoning plays an important role in map construction (Biegler 2000; Clark & Taube 2009; Collett & Collett 2009a,b; Gallistel 1990; McNaughton et al. 2006; Tcheang et al. 2011). It provides the animal with its own approximate

coordinates in a frame that remains the same as the animal moves between widely separated parts of its environment. This makes it possible for the animal to represent in a common system of coordinates the location and orientation of the surfaces it encounters in locations far removed from one another. The mechanism that mediates map construction is the spatial learning mechanism.

Dead reckoning plays an important role in landmark recognition, that is, in the establishment of an identity between a currently perceived distinctive feature of the environment and a charted feature, which is a feature whose location and orientation has previously been recorded on the cognitive map. In natural environments, the readily perceptible features of most landmarks are rarely sufficient in and of themselves to uniquely identify them. The animal’s dead-reckoning-based sense of its location and heading on its cognitive map establishes a prior probability distribution on the charted landmarks that might plausibly correspond to a terrain feature it currently perceives. Landmarks in improbable locations or with an improbable orientation are treated as “impostors” and ignored, no matter how much they may resemble the one that animal is looking to use (Gallistel 1990). Landmarks in the right location and orientation are accepted despite wide variation in their salient surface characteristics, such as color (Cartwright & Collett 1983). Thus, the learning of where you are by dead reckoning is intimately connected to your ability to recognize that what you are looking at now is a unique landmark that you have seen before and represented on your cognitive map.

A stable frame of reference is a *sine qua non* for functional dead reckoning. The summing of successive small displacements (in the limit, the integration of velocity with respect to time) only yields a useable representation of current location if the displacements are all computed within the same frame of reference. If the frame of reference varies during the summation/integration, the resulting vector does not correctly represent the animal’s location in any frame of reference (**Figure 1**). Thus, dead

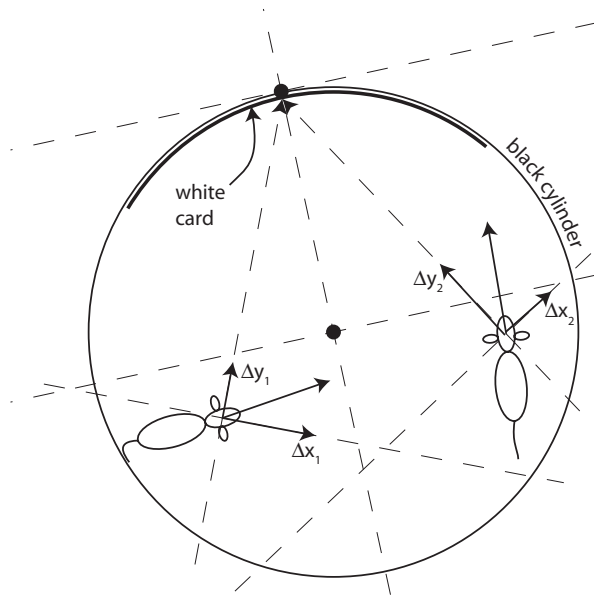


Figure 1

Dead reckoning requires using at every location a frame of reference whose axes are parallel to the frame at the other locations (“carrying the parallel”). It cannot be validly carried out in a changing frame of reference, such as a frame of reference in which the end of one axis is anchored to a prominent landmark. In this example, the landmark is a large white card set against the wall of a black cylinder. If the animal were to use the ever-changing direction from itself to the center of the white card as one axis in its dead-reckoning frame of reference, the vector that results from summing successive displacements, $(\sum \Delta x, \sum \Delta y)$, would not represent its location in any frame of reference. For the resultant vector to be useful, the animal’s displacements in different parts of the environment must be represented in the same coordinate framework and by reference to a single system of directional parallels. In this figure, one such framework has its origin at the center of the cylinder and one axis passing through the center of the white card. This is an array-centered framework whose origin is derived from the geometry of the cylinder. Another has its origin at the center of the white card, with one axis perpendicular to it and the other tangential. This is an object-centered framework, defined by reference to the geometry of the card. Direction for dead-reckoning purposes must be reckoned with respect to what for practical purposes is a point at infinity, a point so far away that its direction does not change as the animal moves. Perceptible terrestrial landmarks are rarely far enough away.

reckoning and map construction are intimately intertwined. Without dead reckoning, map construction is not possible. Without map construction, there is no world-anchored framework within which to represent one’s current position and heading.

Piloting. Piloting is navigation by reference to charted landmarks. It presupposes an

ability to identify currently perceived terrain features with features recorded on a cognitive map. Thus, it presupposes a cognitive map. The map is, of course, learned; the brain is not born with a representation of the geometry of the environment in which the animal happens to find itself. What the brain is born with is the machinery it needs to construct such a representation. This machinery is what enables the animal to learn from experience.

The first and most basic task in piloting is to establish geocentric orientation, orientation within the largest accessible frame of reference. Colloquially, this is called getting one’s bearings. For most animals outdoors, the sun, if visible, is the preferred directional referent. For practical purposes, it is a point at infinity, which means that all lines of sight to it are functionally parallel. Its direction changes because of the earth’s rotation, but animals of diverse species rapidly learn the solar ephemeris, the sun’s direction as a function of the time of day at a given season and latitude, which enables them to compensate for the predictable changes in its azimuth (Dyer & Dickinson 1994, Foa et al. 2009, Gagliardo et al. 2005, Gallistel 1990, Heinze & Reppert 2011, Sauman et al. 2005, Wehner 1984, Wehner & Müller 1993).

Establishing a geocentric orientation without reference to a perceptible point at infinity or to the earth’s magnetic field is an image-registration problem. By “image,” we mean a representation in the brain of a set of perceptible surfaces with substantial relief. Getting oriented under these conditions presupposes two such representations: a cognitive map, constructed from earlier experience in the environment, and a current perception of a portion of the mapped environment, constructed from ongoing sensory input. The map is encoded in one framework. The current perception of the surroundings is encoded in another. The computational challenge is to discover the translation and rotation of the current perception that brings it into register with the corresponding portion of the cognitive map. This computation mediates an animal’s learning which way it is headed when it emerges into

a familiar environment after becoming disoriented with respect to that frame of reference.

There are two basic approaches to image registration: feature matching and computing shape parameters. Feature matching requires finding distinctive features in each image followed by the establishment of cross-image feature correspondences (landmark recognition). It requires that localized regions of an image contain information that makes them unique (hence unambiguously recognizable). Discovering which features in one image “match” which features in another is inherently a trial-and-error process; hence, it is computationally intensive. This contrasts with the closed-form computation of geometric parameters, such as the centroid, principal axes, and higher moments of a shape.

Registration by the computation of shape parameters operates purely on the sets of coordinates that represent the shape. The locations of the surfaces that form a shape, when represented by coordinates, are, from a mathematical perspective, highly structured scatter plots. Image parameters are the centroid, principal axes, medial axes, skews, and so on of these scatter plots; they characterize the spatial distribution of sensible points in the environment by measures computed from the coordinates representing the locations of those points. Image registration by means of shape parameters does not single out features, nor, a fortiori, does it establish between-image feature correspondences; that is, there is no landmark-recognition stage in this computation. The feature-matching approach focuses on local distinctive, easily sensed surface properties, whereas the shape-parameters approach focuses on the global shape of the experienced environment, that is, on its geometry.

The unexpected finding that disoriented(!) rats do not use easily sensed, highly salient corner landmarks to distinguish one end of a rectangular enclosure from the other (Cheng 1986, Margules & Gallistel 1988) led to the suggestion that the image-registration computation mediating the establishment of a geocentric orientation in a disoriented animal

was mediated by the computation of shape parameters (Gallistel 1990). This would explain the rat’s failure to use distinctive nonpositional features to determine which way was which within a rectangular enclosure. If getting reoriented depends on feature matching, the rat’s failure to attend to the landmarks is hard to understand. They were exactly what a feature-matching algorithm most needs—highly distinctive in several sensory modalities and well localized, therefore easily recognized. The hypothesized brain mechanism for establishing orientation on the basis of shape parameters has come to be called the geometric module. There is now a large behavioral literature pro and con (for recent reviews, see Burgess 2008, Cheng 2008, Cheng & Newcombe 2005, Vallortigara 2009). This hypothesis about how a disoriented animal learns its orientation presupposes that the animal has a representation of the overall shape of the experienced environment—a metric cognitive map. On that score, there is now considerable consensus in the behavioral literature, a consensus strongly supported by the to-be-reviewed neuroscientific findings.

Should this hypothesis about the nature of the orientation computation prove correct, landmarks may nonetheless play a role in establishing an orientation. In some environments (e.g., rectangles, which are symmetrical about both principal axes), the principal axes together with the centroid and/or other shape parameters (medial axes, higher-order moments) do not suffice to uniquely orient the navigator. Shape-parameter computations yield two equally acceptable alignments (orientations). Absent other input, the orientation settled on will be wrong half the time. When it is wrong, it will fail to correctly predict salient features (landmarks). This failure may alert the brain to its image-alignment error. It may then try whether the other, equally probably alignment does correctly predict landmark location(s). On this hypothesis, correct alignment precedes and makes possible landmark recognition. On the hypothesis that alignment is achieved by feature matching, alignment follows from landmark recognition. Thus, the question

is not fundamentally about which cues are used, shape parameters, or landmarks; rather, it is about the nature of the image-aligning computation by which an animal becomes oriented on its cognitive map. The nature of the computation determines how the cues are used and the order in which the alignment and recognition processes occur.

The sign-landmark distinction. An important distinction, which has not been clearly maintained in the behavioral literature, and which is rarely recognized in the neurobiological literature, is the distinction between landmarks and signs. A distinctive feature is a landmark when it is used to establish the navigator's bearings (orientation and location in an allocentric frame of reference). A sign, sometimes called a beacon, marks a region where something of motivational interest may be located, something to be approached or avoided. Landmarks cannot be duplicated, because a landmark is, by definition, a unique and recognizable location. In contrast, there can be many instances of a sign. An oak tree, if it is a particularly distinctive one, may function as a landmark, but it more typically functions as a sign that acorns may be found in its vicinity. A storm cloud is a sign that there is bad weather in that direction; it is not a landmark, no matter how salient. The farther away a landmark is, the more effective it is for establishing orientation, whereas the closer a sign is to a goal, the better it serves as a beacon.

There are two considerations of methodological importance for behavioral and neurobiological investigations in connection with the distinction between signs and landmarks: (a) A geocentric reorientation by reference to the shape of, and/or landmarks in, an experimental enclosure is only likely to occur when subjects have been disoriented by prolonged slow rotation in the dark. Absent this inertial disorientation, subjects probably carry a geocentric parallel into an experimental enclosure. In that case, the enclosure is subjectively polarized by its perceived orientation within the broader geocentric framework. Intuitively, the animal that is not geocentrically disoriented

knows which way is which in any enclosure, no matter how symmetrical and how featureless. The violation of this sense of the enclosing maze's geocentric orientation was what confused the rats in the Carr (1917) experiments that first revealed the devastating effect on animal navigation of rotating a maze. (b) The essential test of geocentric reorientation by reference to a putative landmark is the effect of changing that landmark's location on the locus of a subject's search for goals at a substantial distance from the feature. When the goal is at or near or directly behind the distinctive feature and the rat is not disoriented, one is probably not testing properties of the hypothesized geometric module. One is probably testing sign learning (cf. Cheng 2008, Graham et al. 2006, Pearce et al. 2006).

Coordinate formats. Symbol systems for encoding locations differ in how they do it. The most familiar such difference is that between the Cartesian and polar coordinates. In the Cartesian system, the coordinates specify distances from two orthogonal axes. In the polar system, the two coordinates specify an angular deviation from a directional axis (a bearing) and a distance (range). The form of a geometric computation depends strongly on the coordinate format. For computational reasons, it is likely that path integration (dead reckoning, discussed above) is computed in a Cartesian format. In that format, the errors in the estimates of direction are not compounded in the ongoing integration, whereas these unavoidable errors are compounded when the computation is carried on in the polar format. This compounding leads to the rapid buildup of a large error (Cheung & Vickerstaff 2010, Gallistel 1990).

A less familiar means of spatial representation is by spatial basis functions. Spatial basis functions are distributions that may be combined in weighted sums to create a probability distribution that peaks at the subject's probable location. We say more about this less familiar way of representing location when we review the properties of grid cells.

An advantage of the basis-function format is that it naturally encodes spatial probability distributions rather than points. Thus it naturally represents positional uncertainty. Representing positional uncertainty is almost as important as representing position, as many amateur navigators learn to their cost. Also, transformation into a basis-function representation is frequently used in image-registration and image-stitching computations. As already mentioned, computing a geocentric orientation is an image-registration computation. Keeping track of how one local coordinate system relates to the next as one moves through a complex space is closely related to what are called image-stitching problems in image processing. Image registration and image stitching computations map between coordinate frameworks. The maintenance of mappings between different coordinate frameworks is the essence of navigation (cf. Worden 1992).

The place cell system. Neuroscientific evidence for an abstract representation of the geometry of experienced space and for the representation of the animal's location within that geometry (i.e., a cognitive map) comes from the extensive literature on place cells, grid cells, head direction cells, border cells, and boundary-vector cells. As their names suggest, these functionally specialized neurons signal abstract properties of the animal's relation to its spatial environment. These cells are present in rudimentary form as soon as rat pups leave the nest (Langston et al. 2010, Wills et al. 2010), suggesting important genetic control over their development, that is, a genetic basis for the mechanisms by which the brain represents experienced spatial geometry. These specialized neurons are compelling evidence for problem-specific learning mechanisms, that is, mechanisms specialized for learning in a mathematically and physically definable domain of experience. Their specialization for this function makes implicit commitments to domain-defining principles.

Place cells fire when the animal is in a particular place in a familiar environment. In the rat,

where they have been most studied, they are found in the hippocampus, the adjacent subicular complex, and in the entorhinal cortex, which is the main interface between the hippocampus and the neocortex (Moser et al. 2008). Circuits within these three closely connected structures in the medial temporal lobe appear to be specialized for navigational computations.

The sizes of the firing fields for place cells increase as the recording electrode moves from dorsal to ventral within the hippocampus (Jung et al. 1994, Kjelstrup et al. 2008). Viewed from the brain's perspective, different place cells represent spatial location with different degrees of resolution, just as different simple cells in V1 represent local spatial frequencies on the retina with different degrees of resolution.

The striking feature of the firing of place cells is that it does not depend on concurrent sensory input. A place cell fires when a rat stands in or moves through a delimited region of a particular environment. The region is often well away from the walls and other distinctive features of the environment. Different place cells fire in different places. In some environments, a place cell fires regardless of what the rat is looking at, and it fires even if the rat is navigating in complete darkness. Neither immediate visual experience nor prior visual experience is necessary; place cells with normal properties develop in rats blinded soon after birth (Save et al. 1998).

It should be recalled that turning off the lights or blinding rats has little effect on their ability to navigate a familiar maze. The combination of these behavioral and neurobiological findings is not consistent with recurring suggestions in the behavioral literature that places are defined by views (Cheng 2008, Sheynikhovich et al. 2009). Place is defined by reference to coordinates stored in memory that represent the learned geometry of the experienced environment, that is, by reference to a cognitive map. Vision is only one of many different sensory modalities that convey information about the animal's coordinates in a frame of reference. The firing of a place cell represents the integration of and abstraction away from the more sensible

aspects of experience to signal a highly abstract aspect of that experience, namely, its location.

A cognitive map is a repository for acquired information about the positions of sensible features of the animal's environment, as encoded in one or more frames of reference. The firing of place cells is fundamentally dependent on such a repository, which is why the study of such cells is central to the neuroscience of learning and memory.

A striking feature of the results from experiments that have sought to determine what stimuli control the firing of place cells is that removing from the environment a cue that has been shown to affect the place in the experimental environment at which a place cell fires, e.g., a distal landmark, does not terminate its effect on the cell's firing. In many experiments of this kind, rotating a cue by 90° or 180° rotates the cell's place field by a similar amount. However, the cell continues to fire when the animal returns to that place after the cue has been removed altogether (Muller & Kubie 1987, O'Keefe & Speakman 1987, Quirk et al. 1990, Shapiro et al. 1997). What matters is not the current sensory input from the cue, but rather the relation between the rat's current location and the remembered location of the cue.

In another environment, a given place cell will fire to a different place or may not have a field. Thus, the firing of a place cell does not signal that the rat is in a place unique in its experienced world, although the aggregate firing pattern across place cells may do so. In multi-chamber or multiarm environments, a place cell may fire in different places in different chambers or arms (Gothard et al. 1996a,b; Shapiro et al. 1997; Skaggs & McNaughton 1998).

The place where a given neuron fires when in one copy of a box does not predict the box-relative place (if any) where it fires in an exact copy of that box in another room (Leutgeb et al. 2004; O'Keefe & Conway 1978). This reinforces the conclusion drawn from behavioral work that a subject's behavior in one box may differ systematically from its behavior inside a copy of that box when the copies are in different locations in the macro environment, even

when there is no sensory/perceptual access to the macroenvironment from inside the boxes. This re-emphasizes the important point that animals keep track of their own position and orientation in the macroenvironment as well as the position and orientation of the enclosed spaces they enter. They behave differently in different but seemingly identical experimental boxes because the information in memory about the different locations of the two boxes in the macroenvironment informs the brain that the two otherwise indistinguishable boxes are not one and the same (Collett et al. 1997, Collett & Kelber 1988). With enclosures, as with landmarks, location confers identity. No matter how much one enclosure looks and feels and smells like another, it is not that other enclosure if it is not where the animal has the other located on its large-scale map of its experienced environment (or if it has the wrong orientation on that large-scale map). The behavioral results and the neuroscientific results are in accord on this fundamentally important point: Location determines the subjective determination of identity, not vice versa.

Multiple frames of reference. The effects of moving landmarks on the location of place fields suggest conclusions about how the brain uses local landmarks to maintain a geocentric orientation. It does so by using sensed position in a framework established by one or more local landmarks to correct accumulating error in its dead-reckoned orientation. Geocentric orientations based solely on vestibular and optic-flow cues inevitably degrade over time because there is noise in the signals being integrated (Cheung & Vickerstaff 2010, Vickerstaff & Cheung 2010). Frames of reference anchored to perceptible landmarks in an enclosed space can prevent the accumulation of error while the animal is in that environment (**Figure 2**).

Only by keeping track of position in more than one framework can the animal make probable inferences about what is moving with respect to what. It has no way of knowing a priori which objects (potential landmarks) are movable and which are not. If the card moves with

respect to the cylinder, then other points of interest on the circumference of the cylinder (a nest or escape hole, for example) change their coordinates in the card-based frame of reference but not in a cylinder-based frame. To distinguish rotation of the card within the cylinder from rotation of the cylinder itself, the brain must represent places and directions within the cylinder in more than one frame of reference.

When a white card is rotated within an otherwise featureless black cylinder, most place cells change their firing field correspondingly (Yoder et al. 2011). If subjects have been geocentrically disoriented before placement in the cylinder, the card is the only thing that permits the establishment of a stable direction. However, the same shift occurs even when the card is moved while the rat is in the cylinder. This manipulation places the card-centered framework in conflict with cylinder-centered and geocentric frameworks (Blair & Sharp 1996, Sharp et al. 1995). During the small interval over which rotation of the card within the cylinder takes place, there is little accumulation of error in the inertial geocentric framework. If the world (the cylinder, including its floor) were to rotate, there would be a signal from the semi-circular canals. Absent a signal indicating rotation in an inertial frame, the probable inference is that the card moved, not the cylinder. Thus, for locating the nest or escape hole, a coordinate frame anchored to the cylinder is the one to use.

Thus, a fundamental question in these experiments is, when the card is rotated in the presence of a geocentrically oriented subject, does behaviorally measured orientation go with the card? In posing this question, one must not use a behavioral test that puts the goal at the card, because then the card may function as a sign. What is required is a behavioral test of a kind already reported in the literature on animal navigation (Mittelstaedt & Mittelstaedt 1980): A mother gerbil with a nest located behind a hole in a wall of the cylinder leaves the nest to retrieve a pup in the center of the cylinder. If the cylinder (with the nest) is rotated while she gets her pup from a stationary cup in the center, her return run is "correct" in the now erroneous




Figure 2

By maintaining its position concurrently in both local and global frameworks, a navigator can prevent the accumulation of directional error in the global framework. The finely dashed directed line is an initial geocentric parallel; it represents the animal's orientation in a large-scale geocentric framework on entering the enclosed, maximally symmetrical space. The finely dashed circle represents a place field in the framework established by the large white card, which is the only distinctive feature in a black cylinder, an enclosure with no principal axes. In such an enclosure, the only way to maintain a geocentric orientation is by angular dead reckoning, but exclusive reliance on this computation will bring with it unavoidable directional drift (coarsely dashed directed line). The directional drift, ϵ , will put the geocentric framework out of register with the local landmark framework; a position in global coordinates that initially superposed on the place in local coordinates will no longer superpose (coarsely dashed circle). The discrepancy between the referents of what should be corresponding coordinates can be used to correct the drift error in the geocentric direction. Intuitively, if the correct geocentric orientation were that indicated by the coarsely dashed line, then the animal should find itself close to the orientation axis of the card-centered framework. From the quite different gaze angle from the place where a neuron fires, the accumulated error in its geocentric orientation may be computed and corrected.

inertial frame of reference rather than in the cylinder-centered frame. She ignores the highly salient odor cues and pup cries coming from the nest. This is one of many demonstrations of the navigational importance of inertial (hence, geocentric) frames of reference. The question thus is, assuming that a large white card in such a cylinder would establish a frame of reference for place fields, would the rotation of the card

that caused a relocation of the place fields also cause the rat to run 90° off the true direction of a nest located well away from the card? Or, as the hypothesis about the function of place fields in card-centered frames of reference implies, would the rat still return directly to its nest, despite the change in the place fields? To our knowledge, this question has not been put to the experimental test.

When individual landmarks, landmark arrays, or components of more complex environments are moved, the results resist easy summary (Gothard et al. 1996a,b; Shapiro et al. 1997; Tanila et al. 1997), but they are consistent with the hypothesis that the brain's navigation system tracks the animal's position and orientation in several different frameworks simultaneously.

Dependence on vestibular input. Vestibular input is essential to the computational mechanisms that generate the firing of place cells. Temporary inactivation of this signal eliminates the place-specific firing of hippocampal cells for the duration of the inactivation (Stackman et al. 2002). Lesioning the vestibular apparatus eliminates it permanently (Russell et al. 2003). Given the many demonstrations that the place fields depend strongly on local landmarks, the dependence on vestibular input may seem surprising. The explanation is probably that vestibular input is essential to the maintenance of orientation in an inertial (hence geocentric) frame of reference. The maintenance of this orientation is essential to the dead reckoning that plays a critical role in the construction of cognitive maps, landmark recognition, and inferences about what moves relative to what. Whatever disrupts dead reckoning can be expected to disrupt all of these processes.

The evidence that much of the basic machinery of navigation does not operate properly in the absence of appropriate vestibular signals is important for the methodology of behavioral studies of navigation. It suggests caution in interpreting the results from virtual reality experiments and functional magnetic resonance imaging (fMRI) experiments, in which

vestibular signals processed in the midbrain (Angelaki et al. 2010) indicate no translation or rotation in an inertial framework, whereas optic-flow signals processed in the forebrain (Britten 2008) indicate self-motion. Much of the brain's navigational machinery may not function properly under these conditions.

What else place cells signal. Many other aspects of the animal's experience affect the firing of hippocampal place cells. In more complex environments, firing varies strongly with the direction in which the animal moves through a place. Changing the color of the walls has a large effect (Leutgeb et al. 2005), as does changing the task that the animal is carrying out (Colgin et al. 2008, Komorowski et al. 2009, Leutgeb et al. 2005, Manns & Eichenbaum 2009, Markus et al. 1995). Many of these effects cause what is called rate remapping: The place field does not change but the firing rate and firing pattern in that field does. The effects of highly diverse nonspatial cues demonstrate the dependence of neural firing in the hippocampus on a vast repository of acquired information, much of it highly abstract. When we learn how to read the spike train code, the firing of a single hippocampal neuron may tell volumes about the animal's current experience in relation to its past experience in that environment.

It may seem puzzling that a neural structure that is a critical component of a complex system for navigating should also be strongly implicated in memory phenomena that, on their surface, have nothing to do with navigation and the representation of space. Gallistel (1990), in a chapter on "The Unity of Remembered Experience," adduced evidence that spatio-temporal indexing is the mechanism by which the brain knits together the diverse aspects of experience computed by the many different problem-specific modules that are implied by the neuroanatomy and electrophysiology of the cortex. On this hypothesis, the brain binds the remembered color of an object to its remembered shape on the basis that the physically separated (at the level of brain substrates) memories of the object's color and its shape have the same

spatio-temporal index. They have the same spatio-temporal index because they were perceived in the same place at the same time. This hypothesis maintains that episodic memory, that is, the ability to reconstruct an experience in all of its diversity, depends fundamentally on the representation of the spatial and temporal location of the elements of the experience.

Consistent with this hypothesis about the key role in memory of the encoding of spatio-temporal location is the recent discovery that hippocampal place cells also signal position in time. MacDonald et al. (2011) taught rats to associate one of two objects with one of two odors presented 10 s after the rats had inspected the objects. On a given trial, one or the other odor was mixed into the sand in a flowerpot and presented to the rat 10 s after it inspected the object. If the odor was that associated with the recently inspected object, digging in the flowerpot yielded food. If it was the wrong odor, the one associated with the other object, the rat had to avoid digging in the pot and go to another location to obtain food. The experimenters recorded from multiple pyramidal cells in the hippocampus throughout each trial. They found that the firing rates of different cells peaked at different times during the 10 s delay between the inspection of the object and the presentation of the odiferous flowerpot (**Figure 3**). Thus different cells signaled different locations within the interval. As expected from the scalar variability seen in behavioral work on interval timing (Gallistel & Gibbon 2000), the signals from cells whose firing peaked later in the interval were more spread out in time.

The grid cell system. Grid cells are multimodal place cells. A grid cell fires at multiple locations within a familiar environment (see Derdikman & Moser 2010, Moser et al. 2008, Yoder et al. 2011). The locations where it fires form a triangular grid (**Figure 4**). The grids for different cells have different phases, different compass orientations, and different scales, as would be expected if they mediate a basis function representation of the

animal's probable location. The scale of the grids increases as one moves the recording electrode from dorsal to ventral in the entorhinal cortex (Brun et al. 2008). (Recall that a similar increase in the size of place fields occurs along the same dorsal-to-ventral axis in the hippocampus.) The effects of moving landmarks on grid cells are similar to their effects on place and head direction cells; the relevant experiments show that grid cells are anchored to multiple frames of reference and can change their frame of reference within less than 100 ms (Derdikman et al. 2009; Diba & Buzsáki 2008; Frank et al. 2000; Gothard et al. 1996a,b; McNaughton et al. 1996; Redish et al. 2000; Rivard et al. 2004).

An important difference between grid cells and place cells is that the entorhinal grid cell map (together with the head direction cells, border cells, and boundary-vector cells) appears to be a single map with different phase and alignment to the environment in different places or in different conditions (Fyhn et al. 2007, Solstad et al. 2008). The place cell system, by contrast, appears to reflect several different more local maps (Derdikman et al. 2009). An analogy to marine charts may help. Between ports, navigators use a large-scale map on which there are many ports. The navigator uses different parts of this map from different orientations under different circumstances, but it is all one map. Elsewhere in the chart book (or sometimes on the back of the large-scale chart) there are port maps, one for each port. This difference between the grid cell system and the place cell system is consistent with the behavioral evidence that animals maintain their sense of place within the large-scale environment (extramaze cues) while also being sensitive to their location relative to local features (intramaze cues).

Grid cells are found in the medial entorhinal cortex and in the subiculum and parasubiculum, the subset of the medial temporal lobe structures that contain a neural system specialized for the representation of the animal's position and heading on its cognitive map.

Head direction cells fire when the head is oriented within roughly $\pm 45^\circ$ of a




Figure 5

A head direction cell fires whenever the rat's head is at a specific angle with respect to directional parallels (dashed lines), regardless of where the rat is in the environment and even in complete darkness. The experimental set-up portrayed schematizes common features of those actually used: A test arena (radial maze or cylinder or box) surrounded by curtains, often with some landmarks on them (the differently shaped lumps against the curtain), located in a laboratory room off a hall through which the rat is transported prior to testing. Behavioral and electrophysiological results imply that the entire space is represented on at least some of the maps that inform behavior and the firing of place and head direction cells.

directional parallel in some frame of reference (for a recent review, see Taube 2007). The more closely the head's orientation matches the center of a cell's directional tuning, the more rapidly the cell fires. At the optimal orientation, firing is typically brisk (20–100 spikes/s) and sustained. Different cells are tuned to different directions. In the population of head direction cells, there does not appear to be a favored direction. Their directional tuning in a given environment is stable across many days. Thus head direction cells could be described as compass cells. They provide the directional signal required for dead reckoning.

They are not, however, components of a magnetic compass, as their directional tuning is not dependent on the earth's magnetic field.

Head direction cells are found in diverse and widely separated brain structures: the anterior dorsal thalamic nucleus, the lateral dorsal thalamus, and the lateral mammillary nuclei, which are widely separated loci in the diencephalon; also in the dorsal tegmental nucleus in the midbrain; the dorsal striatum in the subcortical telencephalon; and in diverse cortical areas, including entorhinal, retrosplenial, medial precentral, and medial prestriate cortex. They are most prevalent in the anterior dorsal nucleus of the thalamus (~60% of cells recorded there), but the population in the subiculum has also been intensively studied. The subiculum is intermediate between the hippocampus and the entorhinal cortex. Like the entorhinal cortex, it is a way station for signals going into and coming out of the hippocampus. Like place cells, head direction cells require a vestibular signal.

The firing of a head direction cell signals a highly abstract property of the relation between an animal and its surroundings. It does not signal that the head is directed toward a particular place or object in the local environment, because the place toward which the head is oriented differs depending on where the animal is (Figure 5). A head direction cell fires when the head has the cell's preferred orientation even in complete darkness. Direction, like location, is defined only by reference to the learned geometry of the experienced environment.

The effects of moving prominent landmarks on the tuning curves of head direction cells are similar to the effects on place cells: In complex environments, when proximal and distal landmarks are rotated in conflicting directions, the frame of reference usually goes with the distal landmarks (Yoganarasimha et al. 2006, Zugaro et al. 2001), as one would expect, given that for determining direction, the farther away a landmark is, the better it will function. This is another manifestation of the many ways in which the signaling of these specialized cells reveals implicit commitments to domain-specific mathematical principles. In a cylinder

environment, where most such experiments have been done, rotating the card 90° or 180° on the wall of the cylinder rotates the frame of reference for all the head direction cells by the same amount. Notice in **Figure 5** that such a rotation (for example, of the triangle in the northeast quadrant) constitutes a much greater change in viewing angle for the rat when it is looking from the end of the east arm than when it is looking from the end of the south arm. This emphasizes the fact that what rotates coherently is the frame of reference within which directional parallels are defined, not landmark-viewing angles.

When a rat walks from a familiar chamber into an unfamiliar one, a head direction cell typically maintains the frame of reference established in the familiar room (Dudchenko & Zinyuk 2005, Golob & Taube 1999, Stackman et al. 2003, Taube & Burton 1995). This is an example of carrying the directional parallel into unexplored parts of an environment. As the new chamber becomes familiar, the frame of reference mediating the neuron's signaling often shifts to the landmarks in that chamber. However, as the subject gains familiarity with a multichambered environment, some head direction cells adopt a frame of reference that remains the same from chamber to chamber (Dudchenko & Zinyuk 2005).

The shifts in the frame of reference for the head direction signal in response to changes in enclosure shape and landmarks is often taken to indicate that this signal does not participate in the behaviorally well-documented process of maintaining a geocentric orientation with respect to the environment outside the enclosure. This is a mistake. The neuron is not the rat. Carrying a directional parallel depends on computations performed on the overall geometry of the so-far experienced space. Carrying the parallel further, as one enters unexplored regions, is essential to the construction of a coherent map of the large-scale environment. The construction of such a map makes it possible for the place- cum-head-direction-cum-boundary-cell system to signal direction and location in a large-scale framework. It enables

the animal to keep track of where it is in the world.

Border cells fire when the rat is near a compass-oriented boundary (Solstad et al. 2009). What drives the cell's firing is not concurrent sensory input (e.g., a view or feel or touch). What drives firing is a geometric abstraction, the existence of an extended boundary or obstacle to navigation with a particular orientation with respect to the large-scale environment. For example, a boundary cell may fire all along the east side of a north-south wall, whether the lights are on or not. If it is the wall of a square box and the box is elongated parallel to that wall, the firing field now extends all along the elongated wall. When the wall is removed so that the boundary becomes the limit of the navigable surface on which the rat is supported, the cell still fires all along this limit, even though its sensory properties are now radically different. When the rat is moved into other environments with north-south boundaries, the same cell fires all along the east side of those boundaries, too. If a north-south-oriented wall is inserted partway into one of these environments, the cell fires along the east side of that wall *and* along the east side of the enclosing environment. The moving of landmarks and the changing of environmental shape that cause remapping of place and head direction cell firing do not cause remapping of boundary cells; that is, they do not cause them to fire along boundaries with a different compass orientation. The existence of these cells is strong confirmation of the conclusion drawn from behavioral work that rats generally maintain a sense of their orientation in the large-scale environment. Border cells may be special cases of boundary-vector cells.

Boundary-vector cells fire when a limit to navigation lies at some remove in a particular compass direction from the rat, regardless of the color, material, or shape of the boundary and regardless of whether it is a material obstacle to navigation or an immaterial obstacle, that is, the void where the supporting platform ends (Lever et al. 2009). Like border cells, boundary-vector cells do not usually remap in response to the manipulation of landmarks, and they fire

at the same remove and direction from a limit in different environments with different shapes in different locations within the macroenvironment. The longer a boundary cell's vector, that is, the farther from the boundary its firing field is located, the greater is the extent of the field. This suggests scalar uncertainty in the representation of distance as well as time, a result consistent with behavioral results (Durgin et al. 2009).

Border cells and boundary-vector cells are found intermingled with head direction cells and grid cells in what are by now "the usual suspects," that is, the medial entorhinal cortex and the subiculum, which are in the medial temporal lobe next to the hippocampus.

CONCLUSIONS

The mechanisms of synaptic plasticity (e.g., associative LTP or the Hebbian synapse) do not explain the properties of associative learning. The hypothesis that LTP is the mechanism of memory offers no account of how the highly structured, acquired information that mediates animal navigation and the firing of place, head direction, grid, and boundary-vector cells may be carried forward in time in a manner that makes it available to computation on a millisecond time scale. The stored information in the causal chain that informs the firing of cells in the navigation system can change radically in a fraction of a second, as, for example, when the frame of reference for a place or head direction cell changes (Jezek et al. 2011).

The shortcomings of the synaptic plasticity hypothesis highlight the necessity for a more behaviorally and cognitively sophisticated approach to the neuroscience of learning and memory. The literature on the functional properties of neurons in the medial temporal lobe that are sensitive to environmental geometry implies the existence in the brain of genetically specified, purpose-specific computational mechanisms that construct a metric representation of the geometry of the experienced spatial environment and continually signal the animal's location and orientation within that

representation. The neurobiological results testify to the neurophysiological reality of metric cognitive maps, whose existence has been a controversial hypothesis in psychology and cognitive science for decades. Similar conclusions may be drawn from fMRI results in humans (Epstein 2008, 2011; MacEvoy & Epstein 2011; Morgan et al. 2010), but space does not permit a review.

A common feature of space-representing neurons is the highly abstract nature of their tuning. Their tuning cannot be described in terms of the stimuli acting on sensory receptors when the neuron fires. It can only be described by reference to a map in memory and to entities, such as directional parallels and boundaries, which must be constructed by computations performed on that map. A place is not defined by anything acting on the rat's sensorium when it is at that place, and likewise for a head direction. Position and heading are defined only by reference to the learned geometry of the surrounding space. Most of that geometry is not directly sensible by the rat on any particular occasion when it happens to pass through that place on that heading. Under many circumstances, it is dead reckoning that mediates the brain's representation of the animal's location and heading, rather than the processing of contemporary sensory input unique to that place and heading. That is, the animal's representation of itself as being in a certain place with a certain heading depends on computations that took as inputs idiothetic signals generated while it was moving toward that place, rather than signals emanating from extracorporeal features of that place (its feel, its reflectance, the views from that place, the sounds heard there, the odors smelled there, and so on).

The neurobiological evidence confirms the conclusion drawn from studies of animal navigation that dead reckoning is foundational. Dead reckoning is a domain-specific computation that makes an implicit commitment to a domain-specific mathematical principle. It presupposes nontrivial genetically specified neural mechanisms that implement a system of coordinates anchored to a large-scale

frame of reference. The representation of the geometrical relation between a place and the surrounding navigationally important surfaces comes from the cognitive map, which is the repository in memory of the spatial information acquired in past explorations of the environment. Geometric information enters the brain via many different sensory modalities (and, probably, from refferent motor commands) and over extended periods of time.

The extended period of time required for the construction of a global map reminds us that the fundamental function of memory is to carry acquired information forward in time in a computationally accessible form (Gallistel & King 2009). Any hypothesis about the neurobiological mechanism of memory must make clear how the proposed memory mechanism stores structured information and makes it accessible to computation. The hypothesis that synaptic plasticity is the mechanism of memory has yet to meet (or even address) this challenge. In the words of Griffiths et al. (2010, p. 363), "...the single biggest challenge for theoretical neuroscience is not to understand how the brain implements probabilistic inference, but how it represents the structured knowledge over which such inference is defined." The representation of the geometry of the experienced environment is a prime example of structured knowledge. The firing of place- and direction-sensitive neurons gives direct neurobiological evidence that such a representation exists in

neural tissue and that it is every bit as abstract as the term "cognitive map" implies.

The aspects of experience that drive learning in other domains are similarly far removed from elementary sense experience. Associative learning—and the signaling of neurons that participate in it—is driven by contingency, not by the temporal pairing of events. Contingency is a property of the global distributions of events in time, just as shape is a property of the global distribution of surfaces in space. Contingency is comprehended through the computation of entropies, just as shape is comprehended through the computation of locations. The signaling of neurons that participate in reinforcement learning is driven by computations that refer to remembered temporal structure. In both the spatial and the temporal domain, these computations can only be performed on a symbolic representation of where events have occurred—where in space and where in time.

Seen from a broad historical perspective, these conclusions support a materialist form of Kantian rationalism: The brain has genetically specified machinery for the construction of a spatio-temporal probabilistic representation of the experienced world. This machinery is a precondition for what have traditionally been thought of as "elementary" sense experiences, because all remembered experience is localized in space and time, with an explicitly represented degree of uncertainty.

FUTURE ISSUES

1. Where and how is the acquired geometric information that informs the firing of place, grid, head-direction, and boundary-vector cells stored? In the synaptic conductances between neurons in the circuits in which the cells are embedded? In molecules within the cells (e.g., in micro RNAs selected on the basis of their base-pair sequences, or in the methylation patterns on stretches of junk DNA, or in switch-like molecules, of which rhodopsin is an example)? In the conformation of molecules embedded in the synaptic membranes?
2. Where and how is the acquired temporal information that informs the firing of hippocampal cells stored?
3. Is environmental shape encoded using the same code as object shape?

4. Complex computations reduce to sequences of the basic arithmetic operations. How are these operations implemented and at what level of neural structure (circuit, cellular, molecular)?
5. How is it possible for the frame of reference in which a cell signals place or head direction to change in less than 100 ms, given that this frame of reference depends on acquired information stored in memory?

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

The final form of our review has benefitted greatly from comments and suggestions made by far-flung colleagues. For taking the time to make these comments, we offer heartfelt thanks to Ken Cheng, Howard Eichenbaum, Russell Epstein, Rochel Gelman, Edvard Moser, Robert Stackman, Jeffrey Taube, and Barbara Tversky. The views expressed are, of course, those of the authors and not necessarily those of these generous colleagues. C.R. Gallistel's research is supported by NIMH grant RO1MH077027. Louis Matzel's research is supported by grants from the NIA (R01AG029289) and the Office of Naval Research.

LITERATURE CITED

- Abraham WC. 2003. How long will long-term potentiation last? *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 358(1432):735–44
- Andersen RA, Essick GK, Siegel RM. 1985. Encoding of spatial location by posterior parietal neurons. *Science* 230:456–58
- Angelaki DE, Yakusheva TA, Green AM, Dickman JD, Blazquez PM. 2010. Computation of egomotion in the macaque cerebellar vermis. *Cerebellum* 9:174–82
- Antle MC, Silver R. 2005. Orchestrating time: arrangements of the brain circadian clock. *Trends Neurosci.* 28:145–51
- Arcediano F, Escobar M, Miller RR. 2003. Temporal integration and temporal backward associations in humans and nonhuman subjects. *Learn. Behav.* 31:242–56
- Baird AL, Putter JE, Muir JL, Aggleton JP. 2004. On the transience of egocentric working memory: evidence from testing the contribution of limbic brain regions. *Behav. Neurosci.* 118:785–97
- Balsam PD, Drew MR, Gallistel CR. 2010. Time and associative learning. *Comp. Cogn. Behav. Rev.* 5:1–22
- Balsam PD, Fairhurst S, Gallistel CR. 2006. Pavlovian contingencies and temporal information. *J. Exp. Psychol.: Anim. Behav. Proc.* 32:284–94
- Balsam PD, Gallistel CR. 2009. Temporal maps and informativeness in associative learning. *Trends Neurosci.* 32:73–78
- Barnet RC, Grahame NJ, Miller RR. 1996. Temporal encoding as a determinant of inhibitory control. *Learn. Motiv.* 27:73–91
- Beck JM, Latham P, Pouget A. 2011. Marginalization in neural circuits with divisive normalization. *J. Neurosci.* 31:15310–19
- Bermudez MA, Schultz W. 2010. Responses of amygdala neurons to positive reward-predicting stimuli depend on background reward (contingency) rather than stimulus-reward pairing (contiguity). *J. Neurophysiol.* 103:1158–70

- Berridge KC. 2012. From prediction error to incentive salience: mesolimbic computation of reward motivation. *Eur. J. Neurosci.* 35:1124–43
- Biegler R. 2000. Possible uses of path integration in animal navigation. *Anim. Learn. Behav.* 25:257–77
- Blair HT, Sharp PE. 1996. Visual and vestibular influences on head-direction cells in the anterior thalamus of the rat. *Behav. Neurosci.* 110:643–60
- Blaisdell AP, Denniston JC, Miller RR. 1998. Temporal encoding as a determinant of overshadowing. *J. Exp. Psychol.: Anim. Behav. Proc.* 24:72–83
- Bouton ME, Garcia-Gutierrez A. 2006. Intertrial interval as a contextual stimulus. *Behav. Proc.* 71:307–17
- Britten KH. 2008. Mechanism of self-motion perception. *Annu. Rev. Neurosci.* 31:389–410
- Brooks RA. 1991. Intelligence without representation. *Artif. Intell.* 47:139–59
- Brun VH, Solstad T, Kjelstrup KB, Fyhn M, Witter MP, et al. 2008. Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus* 18:1200–12
- Budzynski CA, Bingman VP. 1999. Time-of-day discriminative learning in homing pigeons, *Columbia livia*. *Anim. Learn. Behav.* 27:295–301
- Buhusi CV, Meck WH. 2005. What makes us tick? Functional and neural mechanisms of interval timing. *Nat. Rev. Neurosci.* 6:755–65
- Burger DC, Denniston JC, Miller RR. 2001. Temporal coding in conditioned inhibition: retardation tests. *Anim. Learn. Behav.* 29:281–90
- Burgess N. 2008. Spatial cognition and the brain. *Ann. N. Y. Acad. Sci.* 1124:77–97
- Cajal SR. 1894. La fine structure des centres nerveux. *Proc. R. Soc. Lond.* 55:444–68
- Caporale N, Dan Y. 2008. Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* 31:25–46
- Carr H. 1917. Maze studies with the white rat. *J. Anim. Behav.* 7:259–306
- Cartwright BA, Collett TS. 1983. Landmark learning in bees: experiments and models. *J. Comp. Physiol. A* 151:521–43
- Castro CA, Silbert LH, McNaughton BL, Barnes CA. 1989. Recovery of spatial learning deficits following decay of electrically-induced synaptic enhancement in the hippocampus. *Nature* 342:545–58
- Chater N, Tenenbaum JB, Yuille A. 2006. Probabilistic models of cognition: conceptual foundations. *Trends Cogn. Sci.* 10:287–91
- Chemero A. 2011. *Radical Embodied Cognitive Science*. Cambridge, MA: MIT Press
- Cheng K. 1986. A purely geometric module in the rat's spatial representation. *Cognition* 23:149–78
- Cheng K. 2008. Whither geometry? Troubles of the geometric module. *Trends Cogn. Sci.* 12:355–61
- Cheng K, Newcombe NS. 2005. Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychon. Bull. Rev.* 12:1–23
- Cheng K, Shettleworth SJ, Huttenlocher J, Rieser JJ. 2007. Bayesian integration of spatial information. *Psychol. Bull. Rev.* 133:625–37
- Cheng S, Frank LM. 2011. The structure of networks that produce the transformation from grid cells to place cells. *J. Neurosci.* 197:293–306
- Cheung A, Vickerstaff R. 2010. Finding the way with a noisy brain. *PLoS Comput. Biol.* 6:e1000992
- Cheung A, Zhang S, Stricker C. 2008. Animal navigation: general properties of directed walks. *Biol. Cybern.* 99:197–217
- Chomsky N. 1975. *Reflections on Language*. New York: Pantheon
- Clark B, Taube JS. 2009. Deficits in landmark navigation and path integration after lesions of the interpeduncular nucleus. *Behav. Neurosci.* 123:490–503
- Cole RP, Barnet RC, Miller RR. 1995. Temporal encoding in trace conditioning. *Anim. Learn. Behav.* 23:144–53
- Colgin LL, Moser EI, Moser M-B. 2008. Understanding memory through hippocampal remapping. *Trends Neurosci.* 31:469–77
- Collett M, Collett TS. 2009a. The learning and maintenance of local vectors in desert ant navigation. *J. Exp. Biol.* 212:895–900
- Collett M, Collett TS. 2009b. Local and global navigational coordinate systems in desert ants. *J. Exp. Biol.* 212:901–5

- Collett TS, Collett M. 2004. How do insects represent familiar terrain? *J. Physiol. Paris* 98:259–64
- Collett TS, Fauria K, Dale K, Baron J. 1997. Places and patterns—a study of context learning in honeybees. *J. Comp. Physiol. A* 181:343–53
- Collett TS, Graham P. 2004. Animal navigation: path integration, visual landmarks and cognitive maps. *Curr. Biol.* 14:R475–77
- Collett TS, Kelber A. 1988. The retrieval of visuo-spatial memories by honeybees. *J. Comp. Physiol. Ser. A* 163:145–50
- Crystal JD. 2001. Circadian time perception. *J. Exp. Psychol.: Anim. Behav. Proc.* 27:68–78
- Dayan P, Daw ND. 2008. Decision theory, reinforcement learning, and the brain. *Cogn. Affect. Behav. Neurosci.* 8:429–53
- de Jonge M, Racine RJ. 1985. The effects of repeated induction of long-term potentiation in the dentate gyrus. *Brain Res.* 328:181–85
- Dennis W. 1932. Multiple visual discrimination in the block elevated maze. *J. Comp. Physiol. Psychol.* 13:391–96
- Denniston JC, Blaisdell AP, Miller RR. 2004. Temporal coding in conditioned inhibition: analysis of associative structure of inhibition. *J. Exp. Psychol.: Anim. Behav. Proc.* 30:190–202
- Derdikman D, Moser EI. 2010. A manifold of spatial maps in the brain. *Trends Cogn. Sci.* 14:561–69
- Derdikman D, Whitlock JR, Tsao A, Fyhn M, Hafting T, et al. 2009. Fragmentation of grid cell maps in a multicompartment environment. *Nat. Neurosci.* 12:1325–32
- Diba K, Buzsáki G. 2008. Hippocampal network dynamics constrain the time lag between pyramidal cells across modified environments. *J. Neurosci.* 28:13448–56
- Douglas RJ. 1966. Cues for spontaneous alternation. *J. Comp. Physiol. Psychol.* 62:171–83
- Dudchenko PA, Davidson M. 2002. Rats use a sense of direction to alternate on T-mazes located in adjacent rooms. *Anim. Cogn.* 5:115–18
- Dudchenko PA, Zinyuk LE. 2005. The formation of cognitive maps of adjacent environments: evidence from the head direction cell system. *Behav. Neurosci.* 119:1511–23
- Durgin FH, Akagi M, Gallistel CR, Haiken W. 2009. The precision of human odometry. *Exp. Brain Res.* 193:429–36
- Dyer FC, Dickinson JA. 1994. Development of sun compensation by honeybees: how partially experienced bees estimate the sun's course. *Proc. Natl. Acad. Sci. USA* 91:4471–74
- Edelman GM. 1989. *Neural Darwinism*. New York: Oxford Univ. Press
- Elman JL, Zipser D. 1988. Learning the hidden structure of speech. *J. Acoust. Soc. Am.* 83:1615–26
- Epstein RA. 2008. Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends Cogn. Sci.* 12:388–96
- Epstein RA. 2011. Cognitive neuroscience: scene layout from vision and touch. *Curr. Biol.* 21:R437–38
- Etienne A, Maurer R, Saucy F, Teroni E. 1986. Short-distance homing in the golden hamster after a passive outward journey. *Anim. Behav.* 34:696–715
- Fiorillo CD, Tobler CD, Schultz W. 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299:1898–902
- Flagel SB, Clark JJ, Robinson TE, Mayo L, Czuj A, et al. 2011. A selective role for dopamine in stimulus-reward learning. *Nature* 469:53–57
- Foa A, Basaglia F, Beltrami G, Carnacina M, Moretto E, Bertolucci C. 2009. Orientation of lizards in a Morris water-maze: roles of the sun compass and the parietal eye. *J. Exp. Biol.* 212:2918–24
- Frank LM, Brown EN, Wilson M. 2000. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* 27:169–78
- Fyhn M, Hafting T, Treves A, Moser M-B, Moser EI. 2007. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446:190–94
- Gagliardo A, Vallortigara G, Nardi D, Bingman VP. 2005. A lateralized avian hippocampus: preferential role of the left hippocampal formation in homing pigeon sun compass-based spatial learning. *Eur. J. Neurosci.* 22:2549–59
- Galletti C, Battaglini PP. 1989. Gaze-dependent visual neurons in area V3A of monkey prestriate cortex. *J. Neurosci.* 9:1112–25
- Gallistel CR. 1990. *The Organization of Learning*. Cambridge, MA: Bradford Books/MIT Press. 648 pp.

- Gallistel CR. 1999. The replacement of general-purpose learning models with adaptively specialized learning modules. In *The Cognitive Neurosciences*, ed. MS Gazzaniga, pp. 1179–91. Cambridge, MA: MIT Press. 2nd ed.
- Gallistel CR. 2012. Extinction from a rationalist perspective. *Behav. Proc.* 90:66–88
- Gallistel CR, Gibbon J. 2000. Time, rate, and conditioning. *Psychol. Rev.* 107:289–344
- Gallistel CR, King AP. 2009. *Memory and the Computational Brain: Why Cognitive Science Will Transform Neuroscience*. New York: Wiley/Blackwell
- Gibbon J, Balsam P. 1981. Spreading associations in time. In *Autoshaping and Conditioning Theory*, ed. CM Locurto, HS Terrace, J Gibbon, pp. 219–53. New York: Academic
- Giocomo LM, Moser M-B, Moser EI. 2011. Computational models of grid cells. *Neuron* 71:589–603
- Golob EJ, Taube JS. 1999. Head direction cells in rats with hippocampal or overlying neocortical lesions: evidence for impaired angular path integration. *J. Neurosci.* 19:7198–211
- Gothard KM, Skaggs WE, McNaughton BL. 1996a. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *J. Neurosci.* 16:8027–40
- Gothard KM, Skaggs WE, Moore KM, McNaughton BL. 1996b. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* 16:823–35
- Gottlieb DA. 2008. Is the number of trials a primary determinant of conditioned responding? *J. Exp. Psychol.: Anim. Behav. Proc.* 34:185–201
- Graham M, Good MA, McGregor A, Pearce JM. 2006. Spatial learning based on the shape of the environment is influenced by properties of the objects forming the shapes. *J. Exp. Psychol.: Anim. Behav. Proc.* 32:44–59
- Griffiths TL, Chater N, Kemp C, Perfors A, Tenenbaum JB. 2010. Probabilistic models of cognition: exploring representations and inductive biases. *Trends Cogn. Sci.* 14:357–64
- Groh JM, Trause AS, Underhill AM, Clark KR, Inati S. 2001. Eye position influences auditory responses in primate inferior colliculus. *Neuron* 29:509–18
- Gwinner E. 1996. Circadian and circannual programmes in avian migration. *J. Exp. Biol.* 199:39–48
- Harvey CD, Coen P, Tank DW. 2012. Choice-specific sequences in parietal cortex during virtual-navigation decision task. *Nature* 484:62–68
- Hebb DO. 1949. *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley
- Heeger DJ. 1992. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* 9:181–97
- Heinze S, Reppert SM. 2011. Sun compass integration of skylight cues in migratory monarch butterflies. *Neuron* 69:345–58
- Hull CL. 1930. Knowledge and purpose as habit mechanisms. *Psychol. Rev.* 37:511–25
- Hull CL. 1952. *A Behavior System*. New Haven, CT: Yale Univ. Press
- Jezek K, Henriksen EJ, Treves A, Moser EI, Moser M-B. 2011. Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478:246–49
- Jung MW, Wiener SI, McNaughton BL. 1994. Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. *J. Neurosci.* 14:7347–56
- Kaplan P. 1984. Importance of relative temporal parameters in trace autoshaping: from excitation to inhibition. *J. Exp. Psychol.: Anim. Behav. Proc.* 10:113–26
- Kjelstrup KB, Solstad T, Brun VH, Hafting T, Leutgeb S, et al. 2008. Finite scale of spatial representation in the hippocampus. *Science* 321:140–43
- Komorowski RW, Manns JR, Eichenbaum H. 2009. Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J. Neurosci.* 29:9918–29
- Langston RF, Ainge JA, Couey JJ, Canto CB, Bjerknes TL, et al. 2010. Development of the spatial representation system in the rat. *Science* 328:1576–80
- Lattal KM. 1999. Trial and intertrial durations in Pavlovian conditioning: issues of learning and performance. *J. Exp. Psychol.: Anim. Behav. Proc.* 25:433–50
- Legge ELG, Spetch ML, Cheng K. 2010. Not using the obvious: Desert ants, *Melophorus bagoti*, learn local vectors but not beacons in an arena. *Anim. Cogn.* 13:849–60
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser M-B. 2004. Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305:1295–98

- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser M-B. 2005. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309:619–23
- Lever C, Burton S, Ali Jeeewajee A, O’Keefe J, Burgess N. 2009. Boundary vector cells in the subiculum of the hippocampal formation. *J. Neurosci.* 29:9771–77
- Loomis JM, Klatzky RL, Golledge RG, Philbeck JW. 1998. Human navigation by path integration. In *Wayfinding: Cognitive Mapping and Spatial Behavior*, ed. RG Golledge, pp. 121–51. Baltimore, MD: Johns Hopkins Univ. Press
- MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H. 2011. Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron* 71:737–49
- MacEvoy SP, Epstein RA. 2011. Constructing scenes from objects in human occipitotemporal cortex. *Nat. Neurosci.* 14:1323–29
- Manns J, Eichenbaum H. 2009. A cognitive map for object memory in the hippocampus. *Learn. Mem.* 16:616–24
- Margules J, Gallistel CR. 1988. Heading in the rat: determination by environmental shape. *Anim. Learn. Behav.* 16:404–10
- Markman AB, Dietrich E. 2000. In defense of representation. *Cogn. Psychol.* 40:138–71
- Markus EJ, Qin Y-L, Beonard B, Skaggs WE, McNaughton BL, Barnes CA. 1995. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* 15:7079–94
- Marr D. 1982. *Vision*. San Francisco, CA: Freeman
- Martin SJ, Grimwood PD, Morris RGM. 2000. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annu. Rev. Neurosci.* 23:649–711
- Martin SJ, Morris RGM. 2002. New life in an old idea: the synaptic plasticity and memory hypothesis revisited. *Hippocampus* 12:609–36
- Matzel LD, Held FP, Miller RR. 1988. Information and expression of simultaneous and backward associations: implications for contiguity theory. *Learn. Motiv.* 19:317–44
- Matzel LD, Shors TJ. 2001. Long-term potentiation and associative learning: Can the mechanism subserve the process? In *Neuronal Mechanisms of Memory Formation: Concepts of Long-Term Potentiation and Beyond*, ed. C Holshcer, pp. 294–324. London: Cambridge Univ. Press
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, et al. 1996. Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J. Exp. Biol.* 199:173–85
- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser M-B. 2006. Path integration and the neural basis of the “cognitive map.” *Nat. Rev. Neurosci.* 7:663–78
- Meck WH, ed. 2003. *Functional and Neural Mechanisms of Interval Timing*. New York: CRC
- Menzel R, Greggers U, Smith A, Berger S, Brandt R, et al. 2005. Honey bees navigate according to a map-like spatial memory. *Proc. Natl. Acad. Sci. USA* 102:3040–45
- Merkle T, Wehner JM. 2008. Landmark guidance and vector navigation in outbound desert ants. *J. Exp. Biol.* 211:3370–77
- Miller RR, Kasprow WJ, Schachtman TR. 1986. Retrieval variability: sources and consequences. *Am. J. Psychol.* 99:145–218
- Mittelstaedt ML, Mittelstaedt H. 1980. Homing by path integration in a mammal. *Naturwissenschaften* 67:566–67
- Morgan LK, MacEvoy SP, Aguirre GK, Epstein RA. 2010. Distances between real-world locations are represented in the human hippocampus. *J. Neurosci.* 31:1238–45
- Moser EI, Krippf E, Moser M-B. 2008. Place cells, grid cells, and the brain’s spatial representation system. *Annu. Rev. Neurosci.* 31:69–89
- Muller RU, Kubie JL. 1987. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* 7:1951–68
- Napier RM, Macrae M, Kehoe EJ. 1992. Rapid reacquisition in conditioning of the rabbit’s nictitating membrane response. *J. Exp. Psychol.: Anim. Behav. Proc.* 18:182–92
- O’Keefe J, Speakman A. 1987. Single unit activity in the rat hippocampus during a spatial memory task. *Exp. Brain Res.* 68:1–27
- O’Keefe JO, Conway DH. 1978. Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp. Brain Res.* 31:573–90

- Olsen SR, Bhandawat V, Wilson RI. 2010. Divisive normalization in olfactory population codes. *Neuron* 66:287–99
- Pearce JM, Graham M, Good MA, Jones PM, McGregor A. 2006. Potentiation, overshadowing, and blocking of spatial learning based on the shape of the environment. *J. Exp. Psychol.: Anim. Behav. Proc.* 32:201–14
- Power JM, Thompson LT, Moyer JR, Disterhoft JF. 1997. Enhanced synaptic transmission in CA1 hippocampus after eyeblink conditioning. *J. Neurophysiol.* 78:1184–87
- Quinn WG. 2005. Nematodes learn: Now what? *Nat. Neurosci.* 8:1639–40
- Quirk GJ, Muller RU, Kubie JL. 1990. The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J. Neurosci.* 10:2008–17
- Redish AD, Jensen S, Johnson A, Kurth-Nelson A. 2007. Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychol. Rev.* 114:784–805
- Redish AD, Rosenzweig ES, Bohanick JD, McNaughton BL, Barnes CA. 2000. Dynamics of hippocampal ensemble activity realignment: time versus space. *J. Neurosci.* 20:9298–309
- Rescorla RA. 1967. Pavlovian conditioning and its proper control procedures. *Psychol. Rev.* 74:71–80
- Rescorla RA. 1968. Probability of shock in the presence and absence of CS in fear conditioning. *J. Comp. Physiol. Psychol.* 66:1–5
- Rieke F, Warland D, de Ruyter van Steveninck R, Bialek W. 1997. *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press. 395 pp.
- Rivard B, Lenck-Santini PP, Poucet B, Muller RU. 2004. Representation of objects in space by two classes of hippocampal pyramidal cells. *J. Gen. Physiol.* 124:9–25
- Rocheffort C, Arabo A, André M, Poucet B, Save E, Rondi L. 2011. Cerebellum shapes hippocampal spatial code. *Science* 334:385–89
- Rumelhart DE, McClelland JL, eds. 1986. *Parallel Distributed Processing*. Cambridge, MA: MIT Press
- Russell NA, Horii A, Smith PF, Darlington CL, Bilkey DK. 2003. Long-term effects of permanent vestibular lesions on hippocampal spatial firing. *J. Neurosci.* 23:6490–98
- Sauman I, Briscoe AD, Zhu H, Shi D, Froy O, et al. 2005. Connecting the navigational clock to sun compass input in monarch butterfly brain. *Neuron* 46:457–67
- Savastano HI, Miller RR. 1998. Time as content in Pavlovian conditioning. *Behav. Proc.* 44:147–62
- Save E, Cressant A, Thinus-Blanc C, Poucet B. 1998. Spatial firing of hippocampal place cells in blind rats. *J. Neurosci.* 18:1818–26
- Schultz W. 2006. Behavioral theories and the neurophysiology of reward. *Annu. Rev. Psychol.* 57:87–115
- Shannon CE. 1948. A mathematical theory of communication. *Bell Syst. Tech. J.* 27:379–423, 623–56
- Shapiro ML, Tanila H, Eichenbaum H. 1997. Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. *Hippocampus* 7:624–42
- Sharp PE, Blair HT, Etkin D, Tzanetos DB. 1995. Influences of vestibular and visual motion information on the spatial firing patterns of hippocampal place cells. *J. Neurosci.* 15:173–89
- Shastri L, Ajanagadde V. 1993. From simple associations to systematic reasoning: a connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behav. Brain Sci.* 16:417–94
- Sheynikhovich D, Chavarriaga R, Strössl T, Arleo A, Gerstner W. 2009. Is there a geometric module for spatial orientation? Insights from a rodent navigation model. *Psychol. Rev.* 116:540–66
- Simoncelli EP, Heeger DJ. 1998. A model of neuronal responses in visual area MT. *Vis. Res.* 38:743–61
- Skaggs WE, McNaughton BL. 1998. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *J. Neurosci.* 18:8455–66
- Skinner BF. 1938. *The Behavior of Organisms*. New York: Appleton-Century-Crofts
- Smolensky P. 1986. Information processing in dynamical systems: foundations of harmony theory. In *Parallel Distributed Processing: Foundations*, ed. DE Rumelhart, JL McClelland, pp. 194–281. Cambridge, MA: MIT Press
- Solstad T, Boccara CN, Kropff E, Moser M-B, Moser EI. 2008. Representation of geometric borders in the entorhinal cortex. *Science* 322:1865–68
- Sommer S, Beeren Cv, Wehner R. 2008. Multiroute memories in desert ants. *Proc. Natl. Acad. Sci. USA* 105:317–22

- Stackman RW, Clark AS, Taube JS. 2002. Hippocampal spatial representations require vestibular input. *Hippocampus Online* 12:291–303
- Stackman RW, Golob EJ, Bassett JP, Taube JS. 2003. Passive transport disrupts directional path integration by rat head direction cells. *J. Neurophysiol.* 90:2862–74
- Sutton RS, Barto AG. 1998. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press
- Tanila H, Shapiro ML, Eichenbaum H. 1997. Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus* 7:613–23
- Tanzi E. 1893. I fatti e le induzioni nell'odierna istologia del sistema nervoso. *Riv. Sper. Freniatr. Med. Legale* 19:419–72
- Taube JS. 2007. The head direction signal: origins and sensory-motor integration. *Annu. Rev. Neurosci.* 30:259–88
- Taube JS, Burton HL. 1995. Head direction cell activity monitored in a novel environment and during a cue conflict situation. *J. Neurophysiol.* 74:1953–71
- Tcheang L, Bühlhoff HH, Burgess N. 2011. Visual influence on path integration in darkness indicates a multimodal representation of large-scale space. *Proc. Natl. Acad. Sci. USA* 108:1152–57
- Thompson SM, Mattison HA. 2009. Secret of synapse specificity. *Nature* 458:296–97
- Usherwood PNR. 1993. Memories are made of this. *Trends Neurosci.* 16:427–29
- Vallortigara G. 2009. Animals as natural geometers. In *Cognitive Biology: Evolutionary and Developmental Perspectives on Mind, Brain and Behavior*, ed. L Tommasi, L Nadel, M Peterson, pp. 83–104. Cambridge, MA: MIT Press
- Vickerstaff RJ, Cheung A. 2010. Which coordinate system for modelling path integration? *J. Theor. Biol.* 263:242–61
- Ward RD, Gallistel CR, Jensen G, Richards VL, Fairhurst S, Balsam PD. 2012. Conditioned stimulus informativeness governs conditioned stimulus-unconditioned stimulus associability. *J. Exp. Psychol.: Anim. Behav. Proc.* 38:217–32
- Wehner R, Srinivasan MV. 2003. Path integration in insects. In *The Neurobiology of Spatial Behaviour*, ed. KJ Jeffery, pp. 9–30. London: Oxford Univ. Press
- Wehner R. 1984. Astronavigation in insects. *Annu. Rev. Entomol.* 29:277–98
- Wehner R, Müller M. 1993. How do ants acquire their celestial ephemeris function? *Naturwissenschaften* 80:331–33
- Werner-Reiss U, Kelly KA, Trause AS, Underhill AM, Groh JM. 2003. Eye position affects activity in primary auditory cortex of primates. *Curr. Biol.* 13:554–62
- Wills TJ, Cacucci F, Burgess N, O'Keefe J. 2010. Development of the hippocampal cognitive map in preweanling rats. *Science* 328:1573–76
- Wittlinger M, Wehner R, Wolf H. 2006. The ant odometer: stepping on stilts and stumps. *Science* 312:1965–67
- Worden R. 1992. Navigation by fragment fitting: a theory of hippocampal function. *Hippocampus* 2:165–87
- Wystrach A, Cheng K, Sosa S, Beugnon G. 2011. Geometry, features, and panoramic views: ants in rectangular arenas. *J. Exp. Psychol.: Anim. Behav. Proc.* 37:420–35
- Yoder RM, Clark BJ, Taube JS. 2011. Origins of landmark encoding in the brain. *Trends Neurosci.* 34:561–71
- Yoganarasimha D, Yu X, Knierim JJ. 2006. Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. *J. Neurosci.* 26:622–31
- Zhang J, Berridge KC, Tindell AJ, Smith KS, Aldridge JW. 2009. A neural computational model of incentive salience. *PLoS Comput. Biol.* 5:1–14
- Zhang S, Schwarz S, Pahl M, Shu M, Tautz J. 2006. Honeybee memory: A honeybee knows what to do and when. *J. Exp. Biol.* 209:4420–28
- Zugaro MB, Berthoz A, Wiener SI. 2001. Background, but not foreground, spatial cues are taken as references for head direction responses by rat anterodorsal thalamus neurons. *J. Neurosci.* 21:RC154




Figure 3

The firing of cells in the hippocampus is tuned to location in time as well as location in space. Each row gives the normalized firing pattern from one of more than 30 neurons whose activity was simultaneously recorded on repeated trials during the 10 s delay between object sampling and odor presentation (peak firing indicated by red). The neurons have been ordered from top to bottom in accord with where in the interval their firing peaked. For similar results from neurons in posterior parietal cortex in mice, see Harvey et al. (2012). (Reproduced from figure 2, panel B in MacDonald et al. 2011 by permission of the authors and publisher.)




Figure 4

Schematic rendering of the firing fields of two different grid cells (black and red). The regions where a cell fires are represented by the circles. They are connected by dashed lines to emphasize the triangular structure of the grid. The scale factor for the two grids differs by a factor of 1.5. The arrow points to the unique region where both cells would fire. This shows how the firing field of a place cell could be constructed by thresholding the summed input from these two grid cells (cf. Cheng & Frank 2011, Giocomo et al. 2011). The firing of grid cells represents spatial location in a basis function format.



Contents

Prefatory

Shifting Gears: Seeking New Approaches for Mind/Brain Mechanisms <i>Michael S. Gazzaniga</i>	1
---	---

Biological Bases of Behavior

The Endocannabinoid System and the Brain <i>Raphael Mechoulam and Linda A. Parker</i>	21
--	----

Vision

Synesthesia <i>Jamie Ward</i>	49
--	----

Scene Perception, Event Perception, Object Recognition

Visual Aesthetics and Human Preference <i>Stephen E. Palmer, Karen B. Schloss, and Jonathan Sammartino</i>	77
---	----

Attention and Performance

Detecting Consciousness: A Unique Role for Neuroimaging <i>Adrian M. Owen</i>	109
--	-----

Executive Functions <i>Adele Diamond</i>	135
---	-----

Animal Learning and Behavior

The Neuroscience of Learning: Beyond the Hebbian Synapse <i>C.R. Gallistel and Louis D. Matzel</i>	169
---	-----

Evolutionary Psychology

Evolutionary Psychology: New Perspectives on Cognition and Motivation <i>Leda Cosmides and John Tooby</i>	201
---	-----

Origins of Human Cooperation and Morality <i>Michael Tomasello and Amrisha Vaish</i>	231
---	-----

Language and Communication

- Gesture's Role in Speaking, Learning, and Creating Language
Susan Goldin-Meadow and Martha Wagner Alibali 257

Nonverbal and Verbal Communication

- The Antecedents and Consequences of Human Behavioral Mimicry
Tanya L. Chartrand and Jessica L. Lakin 285

Intergroup Relations, Stigma, Stereotyping, Prejudice, Discrimination

- Sexual Prejudice
Gregory M. Herek and Kevin A. McLemore 309

Social Neuroscience

- A Cultural Neuroscience Approach to the Biosocial Nature
of the Human Brain
*Shibui Han, Georg Northoff, Kai Vogeley, Bruce E. Wexler,
Shinobu Kitayama, and Michael E.W. Varnum* 335

Organizational Climate/Culture

- Organizational Climate and Culture
Benjamin Schneider, Mark G. Ehrhart, and William H. Macey 361

Industrial Psychology/Human Resource Management

- Employee Recruitment
James A. Breugh 389

Learning and Performance in Educational Settings

- Self-Regulated Learning: Beliefs, Techniques, and Illusions
Robert A. Bjork, John Dunlosky, and Nate Kornell 417

Teaching of Subject Matter

- Student Learning: What Has Instruction Got to Do With It?
Hee Seung Lee and John R. Anderson 445

Health Psychology

- Bringing the Laboratory and Clinic to the Community: Mobile
Technologies for Health Promotion and Disease Prevention
Robert M. Kaplan and Arthur A. Stone 471

Research Methodology

- Multivariate Statistical Analyses for Neuroimaging Data
Anthony R. McIntosh and Bratislav Mišić 499

Social Network Analysis: Foundations and Frontiers on Advantage <i>Ronald S. Burt, Martin Kilduff, and Stefano Tasselli</i>	527
--	-----

Indexes

Cumulative Index of Contributing Authors, Volumes 54–64	549
Cumulative Index of Chapter Titles, Volumes 54–64	554

Errata

An online log of corrections to *Annual Review of Psychology* articles may be found at
<http://psych.AnnualReviews.org/errata.shtml>