

Cognitive Critique



REVISITING TOLMAN, HIS THEORIES AND COGNITIVE MAPS

ADAM JOHNSON

*Department of Psychology, Bethel University,
St. Paul, Minnesota*

E-MAIL: adam-johnson@bethel.edu

DAVID A. CROWE

*Department of Biology, Augsburg College,
Minneapolis, Minnesota*

E-MAIL: crowe@augzburg.edu

Accepted 16 November 2008

KEYWORDS

cognitive map, behavior, rodent, Bayes' theorem,
inference, hierarchical learning

ABSTRACT

Edward Tolman's ideas on cognition and animal behavior have broadly influenced modern cognitive science. His principle contribution, the cognitive map, has provided deep insights into how animals represent information about the world and how these representations inform behavior. Although a variety of modern theories share the same title, they represent ideas far different than what Tolman proposed. We revisit Tolman's original treatment of cognitive maps and the behavioral experiments that exemplify cognitive map-based behavior. We consider cognitive maps from the perspective of probabilistic learning and show how they can be treated as a form

hierarchical Bayesian learning. This probabilistic explanation of cognitive maps provides a novel interpretation of Tolman's theories and, in tandem with Tolman's original experimental observations, offers new possibilities for understanding animal cognition.

INTRODUCTION

There is a basic question of why we should return to old papers and ideas about cognition. This question is all the more important when those papers and ideas predate so many conceptual and methodological breakthroughs. The simple answer to our question is that it allows us to reframe our own perspectives on cognition. The descriptions, analogies, reasoning and insights of these original papers and their ideas are unfettered by our modern conceptual and methodological perspectives. While these original papers lack modern precision, they offer a candid view of cognition that does not divide cognition into the neat little boxes and areas that we might today. It is for exactly this reason that we should return to the writings of these astute observers of behavior and puzzle over the connections they make, particularly those that seem most foreign to us, and rediscover the parallels to modern research.

One of the basic questions of cognition is whether cognition can exist in animals other than humans and, if so, what is it and how might it be manifested through behavior? This was the central question within Edward Tolman's work and remains a central issue in subsequent interpretation of Tolman's ideas and our modern discussions of cognition.

The majority of Tolman's work was conducted as psychology struggled to assert itself as a scientific discipline. The struggle to define psychological science resulted in heated discussion on the division between topics that were scientifically tractable and those that were too airy or imprecise for solid scientific treatment, no matter how interesting they might be. Animal cognition was and remains close in proximity to this divide. Tolman's preferred perspective on cognition was based on animal rather than human behavior. He wrote:

There are many findings which suggest the operation of something like a pure cognitive or curiosity need in animals and also findings which indicate the nature of the dependence or the independence of this pure cog-

nitive drive upon such more practical wants as thirst, hunger, sex, fear. Furthermore, we, or at any rate I, see these facts and relationships about cognitive needs more clearly when they have been observed in rats or apes than when they have been merely noted in a common-sense way in human beings. (Tolman 1954, p.536)

In these comments, Tolman distinguishes himself from many of his behaviorist contemporaries who would not treat cognition in either humans or animals. Tolman's perspectives can also be directly contrasted to later views of cognition that emphasized human (and quite often very hyper-intellectualized) aspects of cognition.

For reasons that are not entirely clear, the battle between these two schools of thought (stimulus-response learning versus cognitive learning) has generally been waged at the level of animal behavior.

Edward Tolman, for example, has based his defense of cognitive organization almost entirely on his studies of the behavior of rats — surely one of the least promising areas in which to investigate intellectual accomplishments. (Miller et al. 1960, p.8)

Consequently Tolman's writings, though influential, have often lacked a comfortable place in the history of psychology. A variety of recent investigations of animal behavior have once again emphasized the role of cognition and more recent investigations of human cognition have wondered about correspondences with animal behavior. These recent approaches to cognition parallel Tolman's perspectives on cognition and particularly his ideas relating to cognitive maps.

In the following sections we will revisit Tolman's development of cognitive maps, particularly the underlying features of his earlier work that gave rise to these ideas. Tolman (1948) reviewed a series of five experiments as the basis for cognitive maps. We will explore each of these experiments in turn, citing Tolman's original definitions and their evolution through Tolman's later work with links to more recent investigation. We will consider Tolman's broader views on cognition, his prior assumptions for framing the question of cognition and conclude with a modern (semi-Bayesian) interpretation of his ideas.

TOLMAN AND COGNITIVE MAPS

Tolman developed the idea of a cognitive map as an alternative to the

then-common metaphor of a central office switchboard for learning and memory, typical of stimulus-response formulations. He writes:

We assert that the central office itself is far more like a map control room than it is like an old fashioned telephone exchange. The stimuli, which are allowed in, are not connected by just simple one-to-one switches to the outgoing response. Rather, the incoming impulses are usually worked over and elaborated in the central control room into a tentative, cognitive-like map of the environment. And it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release (Tolman 1948, p.192).

The origins of the cognitive map are evident even in Tolman's early writings (Tolman 1932; Tolman and Krechevsky 1933; Tolman and Brunswik 1935). Tolman's anti-reductionist emphasis on purpose and macroscopic or molar behavior stood in stark contrast to many of the learning theories of his contemporaries. Though Tolman was greatly influenced by Watson's behaviorism and considered himself a behaviorist throughout the majority of his career, his framing of behavioral questions, particularly in relationship to expectancies, signaled a key difference from his contemporaries. This initial position was called purposive behaviorism (Tolman 1932) and provided the foundation for Tolman's later perspectives on animal cognition. Purposive behaviorism can be most succinctly summarized as the idea that animals develop expectancies of their dynamic world and through these expectancies they organize their behavior.

What is expectancy?¹ Tolman struggled with this question throughout his work. His early writings described an expectancy as an animal's multifaceted interaction or commerce with its environment.

For instance, Tinklepaugh (1928) observed that monkeys would engage in "surprise hunting behaviors" when one food type was substituted for another expected food type. When the animal found a reward that differed from its previous experience – even though the reward the animal found was supposedly "just as rewarding" – it would continue to search for a stimulus that matched its previous experience. Such behaviors signaled to Tolman that animals, even those as simple as rats, maintained a set of complex, integrated expectancies of their world (Tolman 1932). Tolman would later call

these complex integrated expectancies cognitive maps (Tolman 1948). His later writings emphasized the use of multiple expectancies to inform behavioral performance, particularly when an animal is faced with a choice (Tolman 1954).

Because the expectancies an animal holds and its use of them develop with experience, different cues or sets of cues that underlie behavior change as a function of experience (Tolman and Brunswik 1935; Tolman 1949). That is, both the content and the use of cognitive maps change with experience. From our modern perspective, this statement and its ramifications may be readily apparent; rats, for instance, tend to navigate using map-like place-learning given certain sets of environmental conditions and experiences whereas they tend to navigate using response-learning given other sets of environmental conditions and experiences (O'Keefe and Nadel 1978; Packard and McGaugh 1996). These ramifications were less evident during Tolman's era and led to one of the more contentious debates on the basis of rat navigation (Hull 1943; Tolman 1948). Tulving and Madigan (1970) describe the debate between place-learning, a position advocated by Tolman, and response-learning, a position advocated by Hull, by stating:

...place-learning organisms, guided by cognitive maps in their head, successfully negotiated obstacle courses to food at Berkeley, while their response-learning counterparts, propelled by habits and drives, performed similar feats at Yale (p. 440).

The debate's resolution has been that animal behavior is dependent on rich sets of cue information during early performance of novel tasks and with overtraining this behavior becomes dependent on an increasingly smaller and more specific set cue information. As such, it is worth noting that Tolman was primarily interested in the case of early, relatively novel learning while many of his contemporaries (such as Hull) were investigating the case of later, overtrained behavior (Restle 1957).

In his original formulation of cognitive maps, Tolman (1948) discussed five basic experiments and their implications to develop his perspectives on cognition:

Latent learning: Learning can occur without observable changes in behavioral performance.² This form of learning can be accomplished with a completely random, largely unmotivated exploration.

Vicarious trial and error: Learning occurs through active in-

vestigation. Vicarious trial and error represents active investigation of the bounds of the signaling stimulus or specific signaled contingencies.

Searching for the stimulus: Learning occurs through active investigation by the animal. Highly salient outcomes yield a search for an environmental change that caused the outcome.

Hypotheses: Development of expectancies require testing and outcome stability. Hypothesis behavior is based on the ordered use of a previous set of behavioral expectancies - that some given change in behavior should produce a change in environmental outcome.

Spatial orientation or short-cut behavior: Behavioral inference is represented by the efficient use of hypothesis behavior in a novel circumstance.

The majority of these ideas have gained a multitude of meanings since Tolman originally presented his formulation of the cognitive map. In order to understand more thoroughly Tolman's conceptualization of cognitive maps and the primary place that expectancy holds within each behavior, we revisit his earlier writings and investigate why he draws on these experiments.

LATENT LEARNING

Latent learning can be defined as learning that occurs in the absence of generally observable changes in behavioral performance. Given appropriate experimental conditions, this learning can be subsequently uncovered. Tolman highlighted latent learning as the first experimental elaboration of cognitive maps to demonstrate that learning occurs across multiple modalities and that learning generally occurs even when that learning is not manifest in behavior. This integrated learning across multiple modalities provides the basis for an animal to perform different sets of behaviors as different needs arise (e.g. thirst, hunger, etc.).

In a traditional example of latent learning (Spence and Lippitt 1946), two sets of fully fed and fully watered animals are allowed to navigate a Y-maze. At the end of one arm of the Y is food and water is available at the end of the other arm. The first set of rats is then water deprived while the second set of rats is food deprived. When the two sets of rats are then placed on the Y-stem, each set runs to the appropriate arm at levels much greater than chance (rats deprived of water run to water and rats deprived of food run to food).

Latent learning provides the fundamental basis for cognitive maps by allowing for organisms to learn covertly. Animals need

not display all that they have learned at a given moment. By beginning with latent learning, Tolman explicitly argues against the assumption that a lack of performance is synonymous with a failure to learn. Performance, Tolman maintains, is not a simple mixture of rewarded responses, but an interaction of need and expectancy. The cognitive map is useful only insofar as it allows an organism to learn and develop a set of expectancies that anticipate potential future needs (in the Y-maze example for instance, the rats did not know whether they would be water or food deprived).

VICARIOUS TRIAL AND ERROR

Vicarious trial and error is a set of experimentally observable behaviors in which an animal attends to and sometimes approaches a specific choice option but does not commit to it (Muenzinger 1938; Tolman 1938, 1939). At choice points within T mazes or within radial arm mazes vicarious trial and error appears as a vacillation between potential options; the rat orients toward one maze arm then reorients toward another until it finally makes its choice.

Although the term vicarious trial and error carries with it a variety of meanings, Tolman appears to have used this description to emphasize the changing interaction an animal has with its environment. Within the context of cognitive maps, Tolman specifically identifies the temporal relationship between vicarious trial and error behavior and task performance as a significant point of interest. In tasks where vicarious trial and error occurs, rats usually display an increase in vicarious trial and error behaviors immediately before dramatic improvements in task performance. Once performance reaches ceiling, vicarious trial and error behaviors diminish. If the task is suddenly changed or the discrimination is made more difficult rats will again display increased vicarious trial and error behavior as they learn the new set of contingencies. Tolman's hypothesis is that vicarious trial and error signals a state of cognitive map reorganization that provides the substrates for changing behavioral performance.

Tolman's description of vicarious trial and error behaviors suggests the simultaneous interaction of two learning processes (Tolman 1938, 1939, 1948). The first learning process allows the animal to parse its environment and to distinguish between and group co-varying stimuli. The second learning process allows the animal to identify environmental reward contingencies related to environmental stimulus information. Failed behavioral performance could

occur due to failure of either learning process: an animal might fail to distinguish a signaling stimulus, it might fail to link the signaling stimulus and reward, or it might fail to identify any reward contingency at all. Within the typical discrimination tasks Tolman used for examining choice behavior in the rat (Tolman 1938, 1939), the interaction of these learning processes was manifest through increased levels of active investigation of a potential signaling stimulus (e.g. a cue card). Because rats displayed increased levels of vicarious trial and error behaviors following changes in either signaling stimuli or reward contingencies, Tolman argued that these behaviors indicated the vicarious sampling of potential outcomes and that such integrated representations formed a critical component of cognitive maps (Tolman 1948).

SEARCHING FOR THE STIMULUS

The set of experiments highlighted by Tolman (1948) describe searching for a stimulus as the post-hoc attribution of an outcome (in most cases a shock) to a stimulus. When the paired stimulus was removed immediately following the outcome, the stimulus attribution was incorrectly made to another stimulus if it was even made at all. This attribution is what has come to be known as the credit assignment problem (Sutton and Barto 1998). Tolman's focus on the active assignment of credit and the temporal order of this assignment (outcome leads stimulus) contrasted with much previous research that focused on the passive transfer of the value or valence of an unconditioned reinforcer to a previously activated stimulus (stimulus leads outcome).

Tolman's perspective is predicated on a set of pre-existing expectancies that an organism has about its environment. It is important to note that the searching for the stimulus described by Tolman (1948) is novel, single-trial learning; that is, an experimental manipulation greatly deviates from the organism's set of pre-existing expectancies. When the organism's expectancies are violated, the organism investigates and assigns credit. While searching for stimulus behaviors are indicative of active credit assignment in single trial learning, an organism will develop another set of active credit assignment behaviors when faced with multiple different problems of the same class. Tolman called these behaviors hypotheses.

HYPOTHESES

In their paper on *The organism and the causal texture of the envi-*

ronment, Tolman and Brunswik (1935) provide an objective definition for hypothesis behavior: the appearance of systematic rather than chance distributions of behavior. The term "objective" denotes the raw statistical nature of the behavior rather than a subjective, putative experience within the organism. Such a statement can be interpreted as either non-sensically trivial or excruciatingly complex. The behavioral performance of an animal that displays hypothesis behavior may not appear to be different than random at a macroscopic level; however, on closer inspection these animals switch from one strategy to another in a rather discrete manner. Citing Krechevsky (1932), Tolman illustrates this point with observations from a sequential Y decision task: for instance, a rat might first attempt all left turns, then all right turns, then a progressive mixture of the two by alternating between right and left turns.

A basic problem with observations of this type is how they should be reported - a point that is particularly salient because different animals rarely follow the same trajectory through a given hypothesis space. At a macroscopic level, animals might select one option at chance probabilities; at a microscopic level, the patterns of behavior might not be predictable. Standard statistical treatments of learning typically describe behavioral patterns in terms of averages. However, such averages across multiple animals often mask incredibly abrupt transitions from chance performance on a task to nearly perfect performance within a single animal (Gallistel et al. 2004). Likely for this specific reason, Tolman resorts to reporting multiple single cases in an attempt to describe the behavior of rats at spatial choice points (Tolman 1938, 1939).³ The basis for hypothesis behavior is functional chunking of the environment.

Tolman and Krechevsky (1933) emphasize that animals do not attempt nonsense intermediate strategies (e.g. attempting to go straight on a T-maze and hitting the wall) but rather shift from one coherent molar pattern of behavior to another. In essence what Tolman has done is allowed the organism to be a good scientist: it maintains a consistent behavior to determine the predictability of the environmental (task) outcome. Tolman's theoretical development suggests that hypotheses are composed of functional chunks, or to use his term, expectancies. A rather simplistic version of hypothesis behavior might be used to explain behavior on the Y-maze during the test portion of latent learning: an approach functional chunk (here a set of motor outputs) combined with the previous learning experience produces a hypothesis, for example given previous experience, approach the deprived substance. The hypothesis

is utilized as a complete chunk and is rewarded. This simple hypothesis behavior is typical of the Y-maze but is easily elaborated to produce much more interesting behaviors within other more complex tasks as we will see.

SPATIAL ORIENTATION

The short-cut behavior in the sunburst maze (Lashley 1929; Tolman et al. 1946) presents an explicit statement of how animals use cognitive maps for problem solving and inference. In this task, rats are pre-trained to navigate through an indirect single sequence of alleys to a food reward site (see Fig. 1 left). During the test phase, the alley sequence from pre-training is blocked and instead multiple radial paths are presented (see Fig. 1 right). One of these radial paths leads directly to the food site. The majority of rats select the radial path that leads directly to food – they select the short-cut (Tolman et al. 1946).




Figure 1 The sunburst maze (after Tolman et al. 1946). Rats were initially trained on the maze shown on the left. Rats were released at the starting point S and rewarded at the maze end R. After training, the rats were tested on the maze shown on the right. The original path out of the circular area was closed and forced the rats to select another arm. Rats primarily selected arms oriented toward the original reward location.

While a number of methodological issues have been identified in this experiment (e.g. a light was placed above the food site in the original experiment; O'Keefe and Nadel 1978), Tolman's argument that "the rats had, it would seem, acquired not merely a strip-map to the effect that the original specifically trained on path led to food but rather, a wider comprehensive map..." (Tolman 1948, p.204) still holds. The short-cut behavior arises from the capacity to combine learned behavioral expectancies with environmental cues

which have not been explicitly rewarded. The appeal of the short-cut has, at its basis, an efficiency argument; the short-cut is the most efficient task solution. The resulting question is over what cue sets and manipulations can an animal continue to produce the most efficient method of task solution. Tolman's descriptions of narrower or wider cognitive maps suggest differences in the depth of information content encoded within a map and, consequently, indicate how informative a map will be over a broad variety of environmental cues and manipulations.

The width of a map might be considered to be a question of specificity and generalization. In contrast to others' subsequent interpretation of cognitive maps (e.g. O'Keefe and Nadel 1978), Tolman (1948) does not indicate how wide or narrow a cognitive map should be or what specific set behaviors are indicative of the cognitive map. Beyond the basic set of illustrative behaviors listed above, Tolman simply states that certain conditions such as brain damage, inadequate environmental cues, over-training, very high motivational or emotional states lead to narrower, strip-like maps behavior. Given appropriately impoverished conditions, cognitive maps can be reduced to very narrow stimulus-response maps; however, the crux of Tolman's formulation of cognitive maps suggests that the deductive approach will not yield the complex interactions made possible through broad cognitive maps and observed in less impoverished environments.

INTERPRETING TOLMAN'S IDEAS

SUBSEQUENT READINGS AND HISTORICAL PERSPECTIVES

An early reading of Tolman by MacCorquodale and Meehl (1954) emphasized the similarity of his theory with other learning theories. The observation that Tolman's use of the term expectancy contained multiple different meanings and implications led MacCorquodale and Meehl (1954) to describe Tolman's formulations as unnecessarily imprecise and, perhaps, even shoddy. Their "precise" reformulation considers Tolman's learning theory in terms of learning to rather than learning that (MacCorquodale and Meehl 1954). The end result of their action-based reformulation was a set of functional and mathematical relations that were virtually indistinguishable from other contemporary theories of learning (c.f. Hull 1943).

The similarity of his reduced theory to Hull's theory and others was not disputed by Tolman (1955) and to a certain extent this perspective was anticipated by Tolman (1948), but neither did Tolman fully agree with this position.

Two critical points identified by MacCorquodale and Meehl (1954) were to become a powerful central influence in subsequent readings of Tolman. These are first that "Tolman's complex cognition statements are not themselves truth-functions of their components"⁴ (p.184) and second that when a rat's expectancies are discussed rather than its habits, the questions of subjective reference and intention become unavoidable. Each point highlights the increasingly prominent position that representation would come to hold within descriptions of animal behavior and particularly within the nascent field of cognitive psychology.⁵ Debate on the representational substrates of learning first centered on whether animals learned place information – the location of reward – or a sequence of responses that led to reward (Hull 1943; Tolman 1948; Restle 1957). These conceptual developments, particularly those related to place representations within cognitive maps, aroused the interest of later cognitive and neurobiological research.

In their influential book entitled *The Hippocampus as a Cognitive Map*, O'Keefe and Nadel (1978) suggested that the then recently discovered place cells found within the hippocampus (O'Keefe and Dostrovsky 1971; O'Keefe 1976) provided the neurobiological substrates of the cognitive map. This assertion revitalized discussion of cognitive maps which had waned with the arguments by MacCorquodale and Meehl (1954) and Restle (1957) that Tolman's learning theory was, at its basis, equivalent to Hull's theory (Hull 1943) and by cognitive psychologists that rat behavior was uninteresting in terms of cognition (see quote by Miller and colleagues above). Subsequent investigations of hippocampal place cell activity further developed and refined ideas about how the hippocampus represents place information and continued to draw connections to cognitive maps particularly beyond spatial domains (Redish 1999). There are, however, several noteworthy conceptual differences between Tolman's original formulation of cognitive maps and the cognitive map developed on the basis of hippocampal place cell research following O'Keefe and Nadel (1978).

One critical development within general theoretical formulations of learning was the identification of multiple memory systems. That some mnemonic functions are devastated following brain area

specific damage while other functions remain intact suggested that memory was not the single unified system previously hypothesized (Scoville and Milner 1957; Cohen and Squire 1980).

Within this context, O'Keefe and Nadel (1978) suggested that cognitive maps mediated by the hippocampus comprise one memory system while another set of extra-hippocampal brain areas subserve a second set of memory function in which learning occurs through processes much like those proposed by Hull (1943). Given only behavioral observations, Tolman (1949) anticipates a number of features of the multiple memory systems debate and provides a rudimentary outline of what he called multiple types of learning.

A second development specific to O'Keefe and Nadel's treatment of cognitive maps was a primary emphasis on the spatial components of place learning in map development. O'Keefe and Nadel (1978) began their arguments for cognitive maps with an epistemological examination of space. The functional parallel drawn by O'Keefe and Nadel on cognitive maps and hippocampal function (and more specifically place cell function) led to a treatment where the cognitive map was elucidated by the spatial behavior of place cells rather than gross animal behavior. The conclusion of such a treatment is that any hippocampal- or cognitive-map- based function must be related to the spatial or place qualities of maps (O'Keefe 1999); as a result, behaviors such as acquisition of trace conditioning (Solomon et al. 1986; Beylin et al. 2001) and retrieval of episodic memories which are also dependent on the hippocampus were considered epiphenomena of the mechanisms of cognitive map function.

While the renewed interest in cognitive maps elicited by O'Keefe and Nadel (1978) cannot be doubted, O'Keefe and Nadel (1978) left a number of issues central to Tolman's formulation of cognitive maps as open questions. O'Keefe and Nadel's perspective on cognitive maps was primarily based on providing a neurobiological basis for latent learning, searching for the stimulus behaviors, and shortcut behaviors (experiments 1,3 and 5). This perspective, however, wholly neglects vicarious trial and error (which goes unmentioned in their text) and only peripherally treats hypotheses (experiments 2 and 4). These topics and a number of other related questions comprise current lines of behavioral and neurobiological research.

The question of how animals switch maps has been and continues to be treated within the place cell literature (Markus et al. 1994, 1995; Touretzky and Redish 1996; Redish 1999; Leutgeb

et al. 2005a,b; Wills et al. 2005). Vicarious trial and error and its hippocampal dependence have begun to receive attention (Brown 1992; Hu and Amsel 1995; Hu et al. 2006; Johnson and Redish 2007).⁶ The mechanisms of path planning and evaluation have also begun to receive increasing attention (Daw et al. 2005; Foster and Wilson 2006; Johnson et al. 2007; Johnson and Redish 2007). It will be interesting to see whether these lines of research will continue to be bound together through cognitive maps or whether they will continue to diverge.

A MODERN COMPUTATIONAL/THEORETICAL/STATISTICAL/ BAYESIAN READING

A systematic formulation of Tolman's learning theory was first approached by MacCorquodale and Meehl (1954). Its reductionist perspective clearly frustrated Tolman (Tolman 1955), but Tolman and the entire field lacked much of the appropriate vocabulary and mathematical and algorithmic sophistication to clearly articulate their ideas. Tolman's treatment of learning seems to fall much in line with recent hierarchical Bayesian approaches to behavior, reasoning and inference (Tenenbaum et al. 2006; Griffiths and Tenenbaum 2007; Kemp et al. 2007). While Tolman did not explicitly formulate his theories in equations, he applauded early attempts to examine probabilistic learning that were the forerunners of modern probability-based (Bayesian) approaches to learning (Tolman and Brunswik 1935).

Hierarchical Bayesian approaches to learning and inference provide a systematic framework that appears to match much of Tolman's conceptual development and, we argue, provide new insights to Tolman's ideas.⁷

Three parallels between Tolman's cognitive maps and hierarchical Bayesian treatments of learning and inference are particularly important to consider. These are (1) the global computations required for both cognitive maps and Bayesian inference, (2) the explicit development of hypotheses for interpreting sparse data, and (3) the changing forms of information utilized to make decisions and how this organizes behavior. We view the frequent discussion of prior distributions as a component of development of hypothesis for interpreting sparse data. Let us provide a brief overview of Bayes' rule (Equation 1) before turning to its use in learning and inference.

$$p(h|o, T) = \frac{p(o|h, T)p(h|T)}{\sum_{h' \in H_T} p(o|h', T)p(h'|T)}$$

Equation 1

At its simplest, Bayes' rule describes the relationship between probabilities $p(a|b)$ and $p(b|a)$ where $p(a|b)$ is read as the probability of a given b . Bayes' rule can be written as where the terms $p(a)$ and $p(b)$ are prior distributions describing the probability of observing a or b (Jaynes 2003). In this form, Bayes' rule forms the basis for statistically appropriate expectations related to variable a based on the coincidence of another variable b . An important computational aspect of Bayes' rule is that the inference $a|b$ requires consideration of every alternative to a . In other words, given the observation of b , we must consider all possible a (that is all $a' \in A$). This global computation is a hallmark of any Bayesian treatment.

More generally, Bayes' rule has been used to examine multiple hypotheses given sparse data. The conditional form of Bayes' rule (Equation 2) has been used to describe the probability of a particular hypothesis h given an observation o and a theoretical framework T .

$$p(a|b) = \frac{p(b|a)p(a)}{\sum_{a' \in A} p(b|a')p(a')} = \frac{p(b|a)p(a)}{p(b)}$$

Equation 2

This conditional form of Bayes' rule suggests a hierarchical framework for inference: competing or overlapping hypotheses can be tested in light of observation data o and a greater, overarching theoretical framework T (see Fig.2). Tenenbaum et al. (2006) have used this framework to show how abstract rules and principles can be used and derived given only sparse language data.

While hierarchical Bayesian perspectives have been used primarily within the arena of human inference (particularly within language), these perspectives can be more broadly applied to behavioral inference of animals, particularly the sort of behaviors and

experiential regimes Tolman used to describe cognitive maps.

In terms of animal behavior, a hierarchical Bayesian treatment suggests that animals form and test hypotheses. These hypotheses structure behavior so that inferences can be made even with relatively little experience (sparse data). In order to develop several important conceptual points related to this treatment, let us examine two sets of behavioral experiments on the partial reinforcement extinction effect (Capaldi 1957) and the use of spatial schemas by rats (Tse et al. 2007).




Figure 2. Hierarchical approaches to Bayes' rule for cognitive research (after Tenenbaum et al. 2006). The left column shows a basic outline for hierarchical Bayesian analysis. Observations are interpreted based on an inferential hierarchy. At the lowest level of this hierarchy are structured probabilistic models that are explicit hypotheses on the distribution of observations. Higher levels allow comparison of multiple probabilistic models relative to data and domain specific structure principles. And these hierarchies can be further extended to include higher order abstract principles. The central column shows how hierarchical Bayesian analysis has been used for taxonomic inference for pictures by Tenenbaum and Xu (2000). Within this example, low hierarchical levels are used for analyzing picture contrast/similarity and higher hierarchical levels are used for category and word selection (taxonomy). The right column shows our interpretation of Tolman's ideas on cognitive inference using a hierarchical Bayesian approach. Following Tolman and Brunswik (1935) and Tolman (1948, 1949), we argue that animals learn the causal texture of the environment and this leads to the formation of cognitive maps and higher order cognitive structure (that provide principles for switching cognitive maps). Hierarchical Bayesian approaches explicitly suggest how cognitive maps fundamentally alter an animal's perception of its environment, its remembrance of prior experience and, consequently, its inference (Tolman 1949).

EXTINCTION AND THE PARTIAL REINFORCEMENT EXTINCTION EFFECT

The partial reinforcement extinction effect (PREE) highlights the importance of prior experience and expectations. The basic observation can be described as follows. While animals initially trained

with deterministic reinforcement schedules (e.g. FR1 lever press schedule) abruptly extinguish responding when the reward is made unavailable, animals that are trained with probabilistic reinforcement schedules display an increased resistance to extinction (Capaldi 1957). Hilgard and Bower (1975) cite Tolman and Brunswik (1935) and their development of an organism's probabilistic interactions with its environment as one of the first theoretical predictions of PREE.

More recent treatments of PREE (Courville et al. 2006; Redish et al. 2007) suggests how the level of uncertainty within the initial reinforcement schedule influences extinction. The initial probabilistic training schedule produces an expectation that a lever press might not yield reward. Such an expectation means that when a lever press does not yield reward, it could be caused by either (1) a continuation of the probabilistic training schedule or (2) a new training schedule. As a result, the expectation or prior developed through the initial training profoundly affects the interpretation of a lever press without reward and how an animal responds. Animals initially trained with a deterministic reward schedule have no basis for interpreting the failed lever press as continuation of the original schedule and can, consequently, interpret the failed lever press as the signal for a new training schedule.

A number of interpretations identify extinction as new learning and not unlearning. These treatments suggest that during extinction a new cognitive state or context is learned wherein a given action fails to deliver on a previously learned expectation (Bouton 2002; Redish et al. 2007). The formation of separate and separable states in which previous learning is valid or is invalid creates an inferential hierarchy based on the construction of new cognitive or inferential structures. Tolman (1949) predicted the formation of such new cognitive states and their influence in inference. He writes:

...in the course of the usual learning experiment there may be acquired not only a specific [*cognitive map*] but also new modes or ways of perceiving, remembering and inferring...which may be then utilized by the given organism in still other later environmental set-ups (italics ours) (p.145).

SPATIAL SCHEMAS

Spatial schemas provide a useful starting point for considering how

new cognitive states contribute to and organize behavior. A recent study by Tse et al. (2007) on spatial schemas in rats highlights several important aspects of cognitive maps and hierarchical Bayesian inference. In this study rats were initially trained on odor-place pairings. The presentation of a specific odor indicated that a reward was available at a specific position within the task arena. One set of rats was trained with six consistent odor-place pairings while another set was trained with six odor-place pairings that varied throughout training. Each group of animals was then trained on two novel odor-place pairings (and these novel pairs were identical between the experimental groups – identical odors and places). The rats trained in the consistent condition learned the two new odor-place pairings while the rats trained in the inconsistent condition did not. Tse et al. (2007) suggest that the rats in the consistent condition were able to form a spatial schema that contributed to later learning while the rats in the inconsistent condition were unable to form this spatial schema.

But what is a spatial schema and how does it organize behavior? The critical distinction between the two conditions in this experiment is the probability that a single odor would signal a single location.

This is what Tolman and Brunswik (1935) called the causal texture of the environment. In the consistent condition, this probability was one while, in the inconsistent condition, the probability was less than one. The argument for schema-use follows that when animals were able to form a spatial schema for the odor-place pairings (consistent condition), they were able to learn the set of novel pairs because they also knew where the locations were not (because they were already taken by other odors). This corresponds to the global computation with Bayes' rule that one must integrate all possible alternatives (all $a' \in A$ and that includes not a) in order to correctly make an inference and precisely match cognitive map function.⁸

Furthermore, the experiment by Tse et al. (2007) shows that animals in the consistent condition were able to learn novel odor-place pairs much more quickly with experience than when they had little experience. That is, animals with increasing experience were able to interpret and structure their behavior appropriately given many fewer training trials. The authors' argument that the development of a spatial schema allows animals to more easily learn fits well within Bayesian treatments of learning from sparse data.

DISCUSSION

The vast majority of Tolman's experiments on animal cognition were focused on understanding inference or, in Tolman's words, the animal's sign-gestalt expectancy. In contrast to the experiments of Hull and his contemporaries, Tolman's experiments examined animal behavior during early learning regimes when animals have relatively little (or sparse) experience. Krechevsky and Tolman's observations that animals appeared to organize and structure their behavior even when many theories of their contemporaries suggested otherwise led them to wonder whether these animals were in fact testing hypotheses and what cognitive factors formed the basis for these hypothesis behaviors (Tolman and Krechevsky 1933). Tolman and Brunswik (1935) suggested that animals maintained some representation of the causal texture of the environment and that this served as the basis for inference. Cognitive maps represent Tolman's best efforts to define how animals learn this causal structure of the environment and, consequently, how this structure organizes behavior.

The difficulties Tolman encountered in examining early learning regimes continue to be a challenge for current cognition research. And while many researchers have sought to avoid this challenge by concerning themselves primarily with asymptotic learning regimes, inference in the face of sparse data remains a central question for cognitive scientists. That Tolman's conceptual development of cognitive maps so well aligns with Bayesian treatments of learning is no doubt a result of continued interest in this question.

There exist three primary points of intersection between Tolman's cognitive maps and Bayesian treatments of learning. The simplest point of intersection is inference based on sparse data. Tolman uses the cognitive map as an explicit attempt to explain the apparent reasoning and structure found within animal behavior, particularly within novel environments and novel experimental perturbations that set one set of cues against another. Short-cuts, hypothesis testing and latent learning highlight this type of inference. This parallels Bayesian treatments of sparse data given a set of competing hypotheses, particularly within studies on cue competition and integration (Battaglia and Schrater 2007; Cheng et al. 2007).

The second point of intersection follows from Tolman's expectancies and "commerce" with environmental stimuli (Tolman 1932). Tolman suggests that an animal will examine its environ-

ment and reduce uncertainty through vicarious trial and error and searching for the stimulus. Each of these types of behavior depend on some representation of the causal texture of the environment and lead to information seeking behaviors when observations and predictions diverge. These integrated representations well match the computations of Bayes' rule: new information propagates throughout the entire network – not simply within a single dimension or information level. And the prediction-observation comparisons highlighted by searching for the stimulus strongly resemble Bayesian filter approaches used in machine learning and robotics (Thrun et al. 2005).

The third point of intersection is the hierarchical approach to learning. Tolman (1949) explicitly identifies multiple interactive types of learning and suggests a hierarchy of learning processes. He argues that these different learning systems have different constraints and their interaction results in “new modes or ways of perceiving, remembering and inferring” (p.145) over the usual course of learning. Hypothesis testing and vicarious trial and error highlight the development of hierarchical inference – Tolman maintains these behaviors represent an active search for higher environmental principles rather than a simple random search for a successful stimulus-response pattern. This hierarchical aspect of learning directly corresponds to Bayesian treatments of language that examine the hierarchical development of rules, principles and theories (Tenenbaum et al. 2006).

Our interpretation of Tolman's learning theory and cognitive maps as a predecessor to hierarchical Bayesian treatments of learning is at odds with previous treatments of cognitive maps. The strong spatial cognitive map position taken by O'Keefe (O'Keefe and Nadel 1978; O'Keefe 1999) holds that the hippocampus represents spatial maps and that anything that is also dependent on the hippocampus occurs through some interaction with space. As noted above, O'Keefe's treatment of cognitive maps is a flattened and reduced form of Tolman's original theory. A recent analysis of Bayesian treatments of animal spatial cognition went so far as to dismiss the strong spatial cognitive map position as untenable (Cheng et al. 2007). Our hope is that our interpretation will reinvigorate discussion of cognitive maps and highlight the comprehensive treatment of Tolman's original formulation. We believe that an integrated approach to the issues of multiple maps, path planning and behaviors such as vicarious trial and error and hypothesis testing will prove valuable for future research.

CONCLUSIONS

The proposal that animals, even those as simple as rats, might possess internal models of the world was not novel within 1930's behaviorist literature (Hull 1930). However, Tolman's treatment of cognitive function in animals left some critics to wonder whether animals would be left "buried in thought".

Guthrie (1935) wrote:

In his concern with what goes on in the rat's mind, Tolman has neglected to predict what the rat will do. So far as the theory is concerned the rat is left buried in thought; if he gets to the food-box at the end that is [the rat's] concern, not the concern of the theory (p.172).

Tolman (1955) saw this type of getting lost in thought as a natural consequence of learning and cognitive function (especially humans and moreover those in academics; Tolman 1954) and he left the problem of how cognition manifests itself in behavior only loosely specified, particularly cognition related to deeper inferential processes. This problem increasingly drove many behaviorists away from cognitive map based formulations of learning and inference.

How cognition is manifest in action continues to be a relevant, open question in research on spatial decision-making, planning and imagery. A recent study by Hassabis et al. (2007) demonstrates that individuals with damage to the hippocampus have profound deficits in spatial imagination. The hallmark of these deficits is not a failure to visualize any particular element of a scene, but an inability to integrate these elements into a coherent whole.

These ideas parallel those of Buckner and Schacter (Buckner and Carroll 2007; Schacter et al. 2007) who suggest that a critical aspect of cognition is the ability to richly imagine potential future circumstances. Daw et al. (2005) and Johnson and Redish (2007) have begun to examine the computational and neural machinery required for these types of cognitive function and the translation of such cognition into behavior.

In conclusion, Tolman's writings have much to contribute to current discussions of cognition. The component processes of cognitive map-based inference parallel current areas of research. Incentive learning and devaluation (Adams and Dickinson 1981; Dickinson 1985; Balleine and Dickinson 1998) recall Tolman's account of latent learning. Planning (Crowe et al. 2005; Balleine and Ostlund 2007; Johnson and Redish 2007) and value-based perceptual re-

organization of the environment (Gallistel et al. 2004; Thrun et al. 2005) parallel vicarious trial and error. Work on the credit assignment problem (Dearden et al. 1998; Nelson 2005) and causal reasoning in rats (Blaisdell et al. 2006) parallel searching for the causal stimulus. Ecological approaches to behavior, particularly foraging (Stephens and Krebs 1987) and non-human problem solving (Emery and Clayton 2004), recall Tolman's descriptions of hypothesis testing. And current research on integrated representation-based inference (Gallistel 1990; Battaglia and Schrater 2007; Cheng et al. 2007; Tse et al. 2007) recall Tolman's descriptions of spatial orientation behaviors. Each of these components of cognitive maps now represent large research literatures by themselves. Tolman's comprehensive treatment of cognition and behavior provides a model for integrating and synthesizing ideas across the vast arena of cognitive science. And as modern cognition research continues to become increasingly specialized, this synthesis will become ever more important.

Tolman's willingness to ask novel experimental questions and reconsider previous experiments, coupled with his childlike delight at tackling seemingly insurmountable complexity (particularly his candid admission of it in print) highlight an important legacy for cognitive science. It challenges current students of cognition to think beyond the limits of current experimental techniques and technologies, across the divisions that have grown between departments and research areas, and to imagine new paths to understanding cognition.

ENDNOTES

¹ It should be noted that questions on expectancies and representation are tightly coupled and neither were adequately addressed by Tolman or his contemporaries. They have continued to cause debate (see Searle 1980 and responses). In a recent study, Schneidman et al. (2003) consider the plastic transformation of visual information to visual representation within very early visual processes. They explain that a monkey sitting in a tree will construct visual representations of a neighboring branch based on a variety of factors including how far the monkey could fall, whether there's fruit on the branch, or whether a predator is near. Consequently, these top-down factors influence even the simplest sensory processes.

² This form of learning should be distinguished from learning latent variables within the context of machine learning.

³ Abrupt transitions in behavioral performance appear to be hallmarks of certain choice tasks, particularly those which elicit vicarious trial and error.

⁴ Put simply, this statement suggests that objective observations about the world are transformed into subjective beliefs. These beliefs are consequently no longer constrained to match observational truth. It should be noted that the origins of critical discussion on an agent's belief are related to top-down versus bottom-up processing that developed in early cognitive psychology.

⁵ It should be noted that while an emerging field of research that would become cognitive psychology was developing in the 1950s, these researchers were primarily physicists, mathematicians and engineers who had become interested in behavior (von Neumann and Morgenstern 1944; Simon 1955; Newell and Simon 1972; Turing 1992) and whose perspectives were relatively independent of Tolman's ideas (Hilgard and Bower 1975).

⁶ These perspectives suggest a larger distributed systems approach in which the hippocampus is one of a number of interactive substrates which mediate cognitive map function (Redish 1999) rather than the hippocampus as the sole mediator of the cognitive map (O'Keefe and Nadel 1978; O'Keefe 1985).

⁷ Even Tolman's light hearted inclination toward the complexity that comes with more holistic treatment of behavior matches Bayesian approaches to behavior (see his final comment in Tolman 1949).

⁸ Tolman argued that both the expectation of food availability at the end of one path and the expectation of food unavailability at the end of other paths contribute to an animal's choice (Tolman and Krechevsky 1933). While this point was derisively noted by MacCorquodale and Meehl (1954), it is well-suited to a Bayesian framework.

ACKNOWLEDGMENTS

We would like to thank A. David Redish and Matthijs van der Meer for discussions and comments on the manuscript. We would also like to thank Bruce Overmier, Celia Wolk Gershenson, Paul Schrater, Dan Kersten and the Cognitive Critique reading group for helpful comments.

REFERENCES

Adams CD, Dickinson A (1981) Instrumental responding following reinforcer devaluation. *Q J Exp Psychol B* 33:109–122

- Balleine BW, Dickinson A (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* 37:407–419
- Balleine BW, Ostlund SB (2007) Still at the choice-point: action selection and initiation in instrumental conditioning. *Ann N Y Acad Sci* 1104:147–171
- Battaglia PW, Schrater PR (2007) Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *J Neurosci* 27:6984–6994
- Beylin AV, Gandhi CC, Wood GE, Talk AC, Matzel LD, Shors TJ (2001) The role of the hippocampus in trace conditioning: temporal discontinuity or task difficulty? *Neurobiol Learn Mem* 76:447–461
- Blaisdell AP, Sawa K, Leising KJ, Waldmann MR (2006) Causal reasoning in rats. *Science* 311:1020–1022
- Bouton ME (2002) Context, ambiguity, and unlearning: sources of relapse after behavioral extinction. *Biol Psychiatry* 52:976–986
- Brown MF (1992) Does a cognitive map guide choices in the radial-arm maze? *J Exp Psychol* 18:56–66
- Buckner RL, Carroll DC (2007) Self-projection and the brain. *Trends Cogn Sci* 11:49–57
- Capaldi EJ (1957) The effect of different amounts of alternating partial reinforcement on resistance to extinction. *Am J Psychol* 70:451–452
- Cheng K, Shettleworth SJ, Huttenlocher J, Rieser JJ (2007) Bayesian integration of spatial information. *Psychol Bull* 133:625–637
- Cohen NJ, Squire LR (1980) Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. *Science* 210:207–210
- Courville AC, Daw ND, Touretzky DS (2006) Bayesian theories of conditioning in a changing world. *Trends Cogn Sci* 10:294–300
- Crowe D, Averbeck B, Chafee M, Georgopoulos A (2005) Dynamics of parietal neural activity during spatial cognitive processing. *Neuron* 47:1393–1413
- Daw ND, Niv Y, Dayan P (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8:1704–1711

- Dearden R, Friedman N, Russel S (1998) Bayesian q-learning. In: Proceedings of the 15th National Conference on Artificial Intelligence (AAAI), pp 761–768
- Dickinson A (1985) Actions and habits: the development of behavioural autonomy. *Philos Trans R Soc Lond B Biol Sci* 308:67–78
- Emery NJ, Clayton NS (2004) The mentality of crows: convergent evolution of intelligence in corvids and apes. *Science* 306:1903–1907
- Foster DJ, Wilson M (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680–683
- Gallistel CR (1990) *The organization of learning*. MIT Press, Cambridge, MA
- Gallistel CR, Fairhurst S, Balsam P (2004) Inaugural article: the learning curve: implications of a quantitative analysis. *Proc Natl Acad Sci USA* 101:13124–13131
- Griffiths TL, Tenenbaum JB (2007) From mere coincidences to meaningful discoveries. *Cognition* 103:180–226
- Guthrie ER (1935) *The psychology of learning*. Harpers, New York
- Hassabis D, Kumaran D, Vann SD, Maguire EA (2007) Patients with hippocampal amnesia cannot imagine new experiences. *Proc Natl Acad Sci USA* 104:1726–1731
- Hilgard ER, Bower GH (1975) *Theories of learning*. Appleton-Century-Crofts, New York
- Hu D, Amsel A (1995) A simple test of the icarious trial-and-error hypothesis of hippocampal function. *Proc Nat Acad Sci, USA* 92:5506–5509
- Hu D, Xu X, Gonzalez-Lima F (2006) Vicarious trial-and-error behavior and hippocampal cytochrome oxidase activity during Y-maze discrimination learning in the rat. *Int J Neurosci* 116:265–280
- Hull CL (1930) Knowledge and purpose as habit mechanisms. *Psychol Rev* 37:511–525
- Hull CL (1943) *Principles of behavior*. Appleton-Century-Crofts, New York

- Jaynes ET (2003) Probability theory. Cambridge
- Johnson A, Jackson J, Redish A (2007) Measuring distributed properties of neural representation beyond the decoding of local variables - implications for cognition. In: Holscher C, Munk (eds) Mechanisms of information processing in the brain: encoding of information in neural populations and networks. Cambridge University Press
- Johnson A, Redish A (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* 27:12176-12189
- Kemp C, Perfors A, Tenenbaum JB (2007) Learning overhypotheses with hierarchical bayesian models. *Dev Sci* 10:307-321
- Krechavsky I (1932) The genesis of "hypotheses" in rats. *Univ Calif Publ Psychol* 6:46
- Lashley K (1929) Brain mechanisms and intelligence: a quantitative study of injuries of the brain. University of Chicago Press, Chicago
- Leutgeb JK, Leutgeb S, Treves A, Meyer R, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005a) Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron* 48:345-358
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005b) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309:619-623
- MacCorquodale K, Meehl PE (1954) Edward C. Tolman. In: Estes W (ed) Modern learning theory. Appleton-Century-Crofts, pp 177-266
- Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs WE (1994) Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus* 4:410-421
- Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA (1995) Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci* 15:7079-7094
- Miller G, Galanter E, Pribram K (1960) Plans and the structure of behavior. Adams Bannister Cox, New York

- Muenzinger KF (1938) Vicarious trial and error at a point of choice: a general survey of its relation to learning efficiency. *J Genet Psychol* 53:75–86
- Nelson JD (2005) Finding useful questions: on bayesian diagnosticity, probability, impact, and information gain. *Psychol Rev* 112:979–999
- Newell A, Simon HA (1972) *Human problem solving*. Prentice-Hall, Englewood Cliffs, NJ
- O’Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurol* 51:78–109
- O’Keefe J (1985) Is consciousness the gateway to the hippocampal cognitive map? A speculative essay on the neural basis of mind. In: Oakley DA (ed) *Brain and mind*. Methuen, London, pp 59–98
- O’Keefe J (1999) Do hippocampal pyramidal cells signal non-spatial as well as spatial information? *Hippocampus* 9:352–365
- O’Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Res* 34:171–175
- O’Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Clarendon Press, Oxford
- Packard MG, McGaugh JL (1996) Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol Learn Mem* 65:65–72
- Redish AD (1999) *Beyond the cognitive map: from place cells to episodic memory*. MIT Press, Cambridge, MA
- Redish AD, Jensen S, Johnson A, Kurth-Nelson Z (2007) Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychol Rev* 114:784–805
- Restle F (1957) Discrimination of cues in mazes: a resolution of the “place-vs-response” question. *Psychol Rev* 64:217–228
- Schacter DL, Addis DR, Buckner RL (2007) Remembering the past to imagine the future: the prospective brain. *Nat Rev Neurosci* 8:657–661

- Schneidman E, Bialek W, Berry MJ (2003) Synergy, redundancy, and independence in population codes. *J Neurosci* 23:11539–11553
- Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatr* 20:11–21
- Searle JR (1980) Minds, brains and programs. *Behav Brain Sci* 3:417–457
- Simon H (1955) A behavioral model of rational choice. *Q J Econ* 69:99–118
- Solomon PR, Schaaf ERV, Thompson RF, Weisz DJ (1986) Hippocampus and trace conditioning of the rabbit's classically conditioned nictitating membrane response. *Behav Neurosci* 100:729–744
- Spence KW, Lippitt R (1946) An experimental test of the sign-gestalt theory of trial and error learning. *J Exp Psychol* 36:491–502
- Stephens DW, Krebs JR (1987) Foraging theory. Princeton
- Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. MIT Press, Cambridge, MA
- Tenenbaum J, Xu F (2000) Word learning as bayesian inference. In: Gleitman L, Joshi A, (eds) 22nd Annu. Conf. Cogn. Sci. Soc. Erlbaum, pp 517-522
- Tenenbaum JB, Griffiths TL, Kemp C (2006) Theory-based bayesian models of inductive learning and reasoning. *Trends Cogn Sci* 10:309–318
- Thrun S, Burgard W, Fox D (2005) Probabilistic Robotics. MIT Press
- Tinklepaugh O (1928) An experimental study of representative factors in monkeys. *J Comp Psychol* 8:197–236
- Tolman E, Brunswik E (1935) The organism and the causal texture of the environment. *Psychol Rev* 42(1):43–77
- Tolman EC (1932) Purposive behavior in animals and men. The Century Co, New York
- Tolman EC (1938) The determiners of behavior at a choice point. *Psychol Rev* 46:318–336
- Tolman EC (1939) Prediction of vicarious trial and error by means of the schematic sowbug. *Psychol Rev* 46:318–336

- Tolman EC (1948) Cognitive maps in rats and men. *Psychol Rev* 55:189–208
- Tolman EC (1949) There is more than one kind of learning. *Psychol Rev* 56:144–155
- Tolman EC (1954) Freedom and the cognitive mind. *Am Psychol* 9:536–538
- Tolman EC (1955) Principles of performance. *Psychol Rev* 62:315–326
- Tolman EC, Krechevsky I (1933) Means-end readiness and hypothesis. a contribution to comparative psychology. *Psychol Rev* 40:60–70
- Tolman EC, Ritchie BF, Kalish D (1946) Studies in spatial learning. I. Orientation and the short-cut. *J Exp Psychol* 36:13–24
- Touretzky DS, Redish AD (1996) A theory of rodent navigation based on interacting representations of space. *Hippocampus* 6:247–270
- Tse D, Langston RF, Kakeyama M, Bethus I, Spooner PA, Wood ER, Witter MP, Morris RGM (2007) Schemas and memory consolidation. *Science* 316:76–82
- Tulving E, Madigan S (1970) Memory and verbal learning. *Annu Rev Psychol* 21:437–484
- Turing A (1992) *Collected works of A.M. Turing: mechanical intelligence*. North Holland
- von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ
- Wills TJ, Lever C, Cacucci F, Burgess N, O’Keefe J (2005) Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873–876

