

Context Learning in the Rodent Hippocampus

Mark C. Fuhs

Mark.Fuhs@cs.cmu.edu

David S. Touretzky

dst@cs.cmu.edu

*Computer Science Department and Center for the Neural Basis of Cognition,
Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.*

We present a Bayesian statistical theory of context learning in the rodent hippocampus. While context is often defined in an experimental setting in relation to specific background cues or task demands, we advance a single, more general notion of context that suffices for a variety of learning phenomena. Specifically, a context is defined as a statistically stationary distribution of experiences, and context learning is defined as the problem of how to form contexts out of groups of experiences that cluster together in time. The challenge of context learning is solving the model selection problem: How many contexts make up the rodent's world? Solving this problem requires balancing two opposing goals: minimize the variability of the distribution of experiences within a context and minimize the likelihood of transitioning between contexts. The theory provides an understanding of why hippocampal place cell remapping sometimes develops gradually over many days of experience and why even consistent landmark differences may need to be relearned after other environmental changes. The theory provides an explanation for progressive performance improvements in serial reversal learning, based on a clear dissociation between the incremental process of context learning and the relatively abrupt context selection process. The impact of partial reinforcement on reversal learning is also addressed. Finally, the theory explains why alternating sequence learning does not consistently result in unique context-dependent sequence representations in hippocampus.

1 Introduction ---

Several theories have posited that the hippocampus is responsible for providing a representation of context and that this representation underpins an animal's performance of a variety of tasks sensitive to hippocampal lesions (Hirsh, 1974; Nadel & Willner, 1980; Jarrard, 1993; Levy, 1996; Wallenstein, Eichenbaum, & Hasselmo, 1998; Redish, 1999, 2001; Hasselmo & Eichenbaum, 2005). These tasks include spatial navigation, sequence

learning, and hippocampal lesion-sensitive conditioning paradigms such as reversal learning. Models have typically focused on one or two of these domains (Levy, 1989, 1996; Schmajuk & DiCarlo, 1992; Gluck & Myers, 1993; Samsonovich & McNaughton, 1997; Wallenstein & Hasselmo, 1998; Redish & Touretzky, 1998; Doboli, Minai, & Best, 2000; Hasselmo, Bodelón, & Wyble, 2002; Hasselmo & Eichenbaum, 2005), developing mechanisms of context learning that are well tailored to the domain of study but do not generalize well across domains.

Moreover, empirical data from both lesion and physiology studies have called into question the viability of existing models. Attractor models of hippocampal place cell remapping (Samsonovich & McNaughton, 1997; Redish & Touretzky, 1998; Tsodyks, 1999; Doboli et al., 2000) cannot account for the gradual separation of maps between similar environments (Jeffery, 2000; Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002). Hippocampal models addressing reversal learning (Schmajuk & DiCarlo, 1992; Gluck & Myers, 1993; Hasselmo et al., 2002) fail to capture the experimental observation that after a series of discrimination reversals, rats can learn to reverse behavior after a single trial (Buytendijk, 1930; Dufort, Guttman, & Kimble, 1954; Pubols, 1962). Models of sequence learning (Levy, 1989, 1996; Wallenstein & Hasselmo, 1998) are challenged by the failure to find sequence-dependent differences in hippocampal place cell activity in some studies (Lenck-Santini, Save, & Poucet, 2001; Hölscher, Jacob, & Mallot, 2004; Bower, Euston, & McNaughton, 2005).

While these previous modeling efforts have offered neural mechanisms to explain how the hippocampus remaps, the work presented here focuses on a related but distinct question: *Why* does the hippocampus remap? More specifically, if remapping serves to create a representation of context, then understanding remapping across problem domains requires a general and concrete definition of what a context is. We propose that context learning may be formalized as decomposing a nonstationary world of hippocampal input patterns (and the sensory input and behaviors they represent) into multiple domains, or contexts, within which the distribution of these input patterns is stationary. What critically defines a context is therefore not a particular class of stimuli (e.g., background cues) or behaviors (e.g., random versus directed foraging), but a set of time windows within which the statistical structure of sensory experiences and behaviors is stable.

Such a definition of context advances its role in prediction: if recent experiences suggest that the present context is *C*, then other prior knowledge about context *C* should also be applicable for the foreseeable future (i.e., until the context changes). The predictive power of a context is determined by both the temporal duration of the context and the (systematic) variability within a context. Contexts that generally last only a short duration fail to be useful in prediction because of the high probability that a different

context with different contingencies will soon become active. Contexts that encompass a highly variable set of sensory experiences and behaviors are also poor predictors insofar as they do not discriminate among the many possibilities within the context.

As a concrete example, consider the serial reversal learning task in which animals must make a choice between two alternatives (e.g., levers, maze arms). During odd blocks of trials, one choice is rewarded; during even blocks, the other is rewarded. As training progresses, the two reward contingencies may be codified as two different contexts, each with a static reward structure in which only choice 1 is rewarded (context 1) or only choice 2 is rewarded (context 2). Thus, when a particular choice is rewarded (or not rewarded), it suggests which of the two contexts is active, thereby predicting that that choice will be rewarded on subsequent trials as well. If, however, both reward contingencies were grouped together in the same context, the identity of the active context would not be useful in predicting an optimal behavioral strategy. By contrast, in a random foraging paradigm, the random scattering of food within the environment ensures that the reward location is unpredictable. Dissecting a random foraging session into a large number of contexts, each representing a distinct location at which food was found, would result in a complex contextual representation with no predictive value.

Inferring the most informative context model is therefore the central problem of interest. In this article, we formulate context learning as a model selection problem: Into how many contexts should the world be divided? We develop a statistical framework for context learning based on the hypothesis that the degree of remapping in hippocampus reflects two independent factors: the degree of similarity between contexts and the animal's confidence that the two contexts are in fact distinct. The framework shows how contextual representations should evolve over time as the animal progresses through a training regimen, whether that regimen involves different environments, reward contingencies, or sequences. Within the framework, model selection is biased to prefer contexts that are active for longer periods of time, a key constraint that explains the development of hippocampal contextual representations in spatial and reversal learning paradigms and justifies their absence in overlapping sequence learning. The framework also distinguishes between context learning (inferring the context model), which may be a gradual process over many blocks of trials, and context selection (inferring the current context given a context model), which should be a more abrupt phenomenon. This distinction is critical to understanding serial reversal learning: substantial training is required for rats to achieve single-trial reversals, but once they are trained, single trial reversals can be realized as a context switch. Several testable predictions are made to motivate future experimental work.

2 A Statistical Framework for Context Learning

2.1 Overview

2.1.1 Hidden Markov Models and Hippocampal Activity Patterns. We model context learning as a process in which the hippocampus constructs a generative model of its inputs. Within our statistical framework, which is based on hidden Markov models (HMMs), generative models are composed of states and contexts. An HMM with a given number of states is parameterized by a distribution of expected values for each state and a transition matrix, which expresses the probability of transition from one state to another. Each state represents a conjunctive encoding of the hippocampal input and may be thought of as a particular hippocampal activity pattern. Different states are therefore used to represent different positions within an environment or different stages in a task.

A context is simply a group of states. States are grouped together into contexts such that state transitions are frequent but transitions between states in different contexts are rare. Transition probabilities between states within the same context are not constrained, and each is individually parameterized, but the transition probability between states in different contexts is fixed at a small value. If, for a particular model, context switches occur frequently, the goodness of fit of the model will therefore be judged to be low. Thus, the fixed intercontext transition probability encourages context switches to be rare, or, inversely, context durations to be long.

What is the relationship between states, contexts, and hippocampal activity patterns? In an HMM, states are “identifiable” in the sense that there is an abstract label for each state that is independent of its particular parameters. This permits an HMM to contain multiple distinct states with the same parameters. This state identifiability is a theoretical construct that we view as unreasonable to apply to the hippocampus. Rather, we argue for weaker identifiability constraints:

- States within a context are distinguished only by their expected observations. Thus, within a context, different hippocampal activity patterns can be observed only when different input patterns are presented.
- Contexts are identifiable—the hippocampus forms a latent representation of context. This implies that the same input pattern presented in two different contexts will result in different hippocampal activity patterns.

These constraints suggest that an observed hippocampal activity pattern (i.e., the currently active state) is determined by two factors: afferent activity and the currently active context. Intuitively, when the same input pattern is represented by multiple hippocampal codes, each code must be part of a different context. These constraints are closely related to the behavior of

latent attractors (Doboli et al., 2000), where network activity is determined by both the input pattern driving the network and the network's current attractor basin (which represents the current context).

In order to handle input pattern noise, each state is parameterized by a distribution of input patterns. Thus, determining which state is active is a (minor) statistical inference problem, requiring the determination of the state under which the input pattern is most likely. This is qualitatively similar to pattern completion in associative memory models, where the network activity pattern resulting from a noisy input is made more similar to the best matching previously stored pattern. Several theories have implicated the hippocampus in conjunctive encoding (e.g., Rudy & Sutherland, 1995) and associative memory function (Marr, 1971; O'Reilly & McClelland, 1994; Rolls, 1996), and the formulation presented here is not incompatible with these notions. However, our work focuses not on the utility of the individual states but on their contextual organization.

2.1.2 Adopting New Models. In the hippocampus, new experimental conditions are observed to cause the creation of new place cell maps. In the model, a new experimental condition is represented by augmenting the current HMM with new states in a new model context. Given that the input patterns are noisy, under what conditions should a new context be added to the model? A larger, more complex model, one with more contexts, will inevitably provide a better fit of the input patterns. However, one should distinguish the extent to which the model is fitting the statistical structure of the inputs from the extent it is better fitting the noise.

In a Bayesian setting, a model should be adopted when its posterior likelihood is higher than that of competing models. This posterior probability accounts for the complexity of the model by averaging the goodness of fit of the model over the prior uncertainty in the model's parameterization. Larger models have more parameter values that must be specified, so the prior probability of any particular parameterization is lower. Larger models are therefore implicitly penalized for the size of their parameter space.

In clear cases, such as when entering a completely new and different environment, the sensory input would differ strongly from what is expected by any state under the "current" model. A larger model, augmented with a group of states for the new environment, would therefore be immediately justifiable. In more subtle cases, such as when a new environment is similar to a familiar environment, the posterior likelihood of the larger model would rise more gradually as the animal gains more experience in each environment. This increasing experience offsets the penalty caused by the larger model's added complexity.

A distributed neural representation of context allows multiple models to be expressed simultaneously. Returning to the two-environment example, if a larger model were only weakly favored over a smaller one, then it is possible for some cells to represent the larger model, distinguishing between

the two environments, while other cells represent the smaller model, showing the same firing patterns in both. We interpret gradual increases in the number of cells that remap between contexts (gradual remapping) as observed by Lever et al. (2002) and others as a reflection of the relative degree of acceptance of each model: the increasing degree of remapping between two environments reflects the increased statistical likelihood of a model that represents them as two separate contexts.

2.1.3 Independent and Dependent Contexts. Consider again the case of two similar environments, where the input patterns at corresponding positions in each environment are similar. These two environments could be represented by a one-context HMM, where each of the p states represents a corresponding position in the two environments. Each state would therefore be optimally parameterized when tuned to the distribution of input patterns observed at that position in either environment. The two environments could also be represented by a two-context HMM, where two different groups of states, each of size p , represent the two environments. Such a two-context model would have double the number of parameters, and the Schwarz criterion (Schwarz, 1978) suggests that this corresponds to a quadratic increase in complexity.

Penalized for such high complexity, preliminary simulations showed that, compared to a one-context model, a two-context model would be considered astronomically improbable until after substantial experience in both environments. This would predict that rats without substantial experience should never show any remapping between similar environments, a finding incompatible with the observation of gradual remapping.

Instead of asking whether there is sufficient evidence to justify the complexity of an entirely new context, one might ask whether there is sufficient evidence to justify the complexity required to express how the second environment differs from the first. If the two environments are similar, then it may be simpler to express this difference than to express the second environment independently. This is tantamount to assuming a priori that the two contexts are related but that their differences may nonetheless be valuable to represent in separate contexts.

A two-level contextual hierarchy is therefore considered in which contexts may be either independent or dependent. At the top level are independent contexts, whose states' parameters are not statistically associated with those in any other independent context. The creation of an independent context must be justified with respect to the complexity of its entire set of parameters, which is easily done when the animal enters an environment clearly unrelated to any other. At the second level are dependent contexts, each associated with an independent context. Specifically, each state in a dependent context is paired with a corresponding state in an independent context, and states in the dependent context may share parameters with states in that independent context. In addition, the expected hippocampal

inputs of states in the dependent contexts are assumed to be similar to those of the states in their respective independent contexts. This reduces the added complexity of the dependent context to more accurately reflect the degree to which it differs from the independent context.

There is a growing body of evidence to support the notion that similar contexts are not represented independently *in vivo*. Very similar contexts show partially overlapping representations (partial remapping), especially when only a change in task is involved (e.g., Markus et al., 1995; Shapiro, Tanila, & Eichenbaum, 1997; Skaggs & McNaughton, 1998; Wood, Dudchenko, Robitsek, & Eichenbaum, 2000; Jeffery, 2000). Even between contexts with more pronounced differences (arena wall shape or color changes), the type of remapping observed primarily involves a change in firing rate, but not of the location of the place field (S. Leutgeb et al., 2005). We therefore interpret the degree of remapping to be determined, in the asymptote, by the similarity of the contexts. With little experience, the degree of remapping is even less, reflecting the uncertainty that the two contexts are distinct. We do not propose a specific similarity metric—any such metric would vary among subjects—but such a metric should be monotonic: changes that make two contexts more different should not result in less remapping.

Even if the contexts are not independent, why should one be represented as an independent context and the other as a dependent context? (In other words, why not represent them symmetrically as two codependent contexts without introducing an explicit hierarchy?) The asymmetry allows the models to be nested, which provides a more direct relationship between model contexts and place cell activity patterns. Consider an experiment involving two similar environmental contexts (E_1 and E_2). Before the start of the experiment, the animal has contextual representations only for its home cage, a transport box, or something similar; let this be context model M_0 . Upon introduction to E_1 , which is nothing like any previously experienced context, a new model is immediately adopted, $M_1 = M_0 \cup \{I_1\}$, that is, all contexts in M_0 and an independent context I_1 for E_1 . Subsequent exposure to E_2 leads to consideration of $M_2 = M_1 \cup \{D_2^{I_1}\} = M_0 \cup \{I_1, D_2^{I_1}\}$, where $D_2^{I_1}$ is a dependent context associated with I_1 that represents E_2 . As more experience is acquired in E_1 and E_2 , the relative likelihood of M_2 will gradually increase. Given a nested model structure, the only change from M_1 to M_2 is the creation of a new context $D_2^{I_1}$ for E_2 . The model context I_1 for E_1 exists (has been statistically justified) independent of any experiences in E_2 . Thus, the most parsimonious hippocampal representational change would involve the creation of a contextual representation for E_2 without changing the representations of any other contexts.

Gradual remapping appears to show such an asymmetry. Lever et al. (2002) first exposed rats to a cylindrical arena (E_1) and then a square arena (E_2). Over the course of many subsequent exposures to both arenas, rats gradually remapped between them. The observed pattern of remapping

suggests that the activity of place cells in the square arena (E_2) diverged from the cylindrical arena E_1 ; in contrast, fields in the cylindrical arena (E_1) appeared stable over several days.

It should be noted that whereas the experiments and simulations in this study focus on learning to distinguish two contexts, the statistical framework outlined here can be extended in a straightforward manner to allow an arbitrary number of dependent contexts.

In the following sections, a detailed description is provided of how simulated hippocampal inputs are constructed, how HMMs with independent and dependent contexts are defined, and how the posterior likelihood of different models is calculated.

2.2 Simulated Hippocampal Inputs. While the inputs to the hippocampus are very high dimensional, input patterns to the model, for computational simplicity, were formulated as noisy scalar values. Input patterns were generated based on a simulated “ground truth” of the animal’s actual state (e.g., location, task stage) in the world. These environmental states (e-states) compose the true generative model, not to be confused with states in the hippocampal HMM model. Each e-state was assigned a positive index i , and the hippocampal input value y generated for that e-state was $8(i - n/2) + \eta$, where n is the number of e-states in the environment and η is a gaussian noise term; the y values are therefore roughly centered around 0, as expected under the prior (see below). The standard deviation of the noise term was 0.125, which ensures that different e-states within an environmental context are unambiguously distinct.

To simulate an experiment involving multiple sequences or tasks, a single set of e-states was used; only their order of presentation changed. For example, multiple spatial sequences were simulated using the same set of position e-states, but the order of visited positions was different for each sequence. To simulate small environmental changes, each e-state’s input value was mildly perturbed (see specific simulations for details).

Experiments were modeled in discrete time: each discrete time step had an associated e-state. To model a specific experimental paradigm of duration T , a sequence of T e-states was produced, representing the entire course of the experiment over many days. A sequence of hippocampal inputs, y_1, \dots, y_T , was then generated based on the e-state sequence.

2.3 Hidden Markov Models with Independent and Dependent Contexts. Our simulations used gaussian HMMs as a model of the hippocampal representation of its input patterns. By using HMMs, we do not intend to suggest that the hippocampus is a finite state machine; rather, HMMs provide a convenient statistical framework in which to represent both a mixture of many different input patterns and their temporal structure (or a Markov approximation thereof).

Formally, a gaussian HMM is composed of a set of N_S states, each parameterized by a normal distribution of expected values, $N(\mu_s, \sigma_s^2)$. A transition probability matrix, A_{ji} , defines the probability of transitioning from state i to state j at each time step. Additionally, a starting state probability p_0 must be specified; in these simulations, the starting state probability was assumed to be uniform over all states.

In our framework, states are organized into multiple contexts, an organization realized by restrictions placed on transition probabilities. Transition probabilities between states in the same context could vary to fit the observed sequence of transitions in y . However, the sum of transition probabilities to all states in other contexts was fixed at $\gamma = 0.05$. This is the critical parameter that determines how strongly models are biased against context changes.

If the total between-context transition probability is γ , what is the transition probability from some state s to some other state s' in a different context? This is defined based on a hierarchical organization of states and contexts. Consider a model M_k composed of N_S states, S_1, \dots, S_{N_S} . These states are organized into N_C contexts, C_1, \dots, C_{N_C} . Contexts are organized into N_G context groups, G_1, \dots, G_{N_G} , where each context group contains exactly one independent context and any dependent contexts associated with it. For the purpose of defining transition probabilities, the model M_k is therefore organized as a mixture of context groups, each context group a mixture of contexts, each context a mixture of states. For each mixture, the component probabilities are set to be equal. Thus, the probability of transitioning to a particular context group, G_g , is $p_1 = 1/N_G$. The probability of transitioning to a particular context in that group is $p_2 = 1/|G_g|$, the inverse of the number of contexts in group G_g . The probability of transitioning to a particular state in that context is $p_3 = 1/|C_c|$, the inverse of the number of states in context C_c . Thus, the probability of transitioning from state s to state s' is $\gamma p_1 p_2 p_3$. When the values p_1 and p_2 are calculated, the context of s is first excluded, since these transitions are only between states in different contexts.

The hierarchical organization of contexts helps us take into account the fact that there are other unrelated contexts outside the experimental apparatuses. A fixed value of N_G is used to reflect that the transition probabilities among unrelated contexts would not change with the addition of one or two more. For the multiple environments simulations, $N_G = 3$, since up to three independent contexts are considered; for the other simulations, $N_G = 1$. (How exposure to just a few or to many completely different environments affects remapping is not known, but an increased willingness to remap after experiencing many previous environments would argue instead that N_G is explicitly represented in the hippocampus.)

We now describe how states in independent and dependent contexts are parameterized. In particular, we describe a model in which a state in a dependent context may share parameters with a state in an independent

context, and the prior probabilities of parameters in dependent contexts may depend on corresponding parameters in their respective independent contexts. Qualitatively, the particular scheme for parameter sharing and the priors used here have a singular purpose: to reduce the size of the parameter space (i.e., the complexity) of the dependent context. Other formulations that achieve the same goal would have been equally acceptable, so readers not well versed in statistics may safely skip the remainder of this section.

To define the parameters of a model, consider a state \hat{s} in a dependent context and its corresponding state s in the associated independent context. The independent state's parameters are defined purely in terms of the state's transition probability vector, \vec{a}_s , and mean and variance terms, μ_s and σ_s^2 :

$$A_{*\hat{s}} = \vec{a}_s \quad (2.1)$$

$$p(y|s) \sim N(\mu_s, \sigma_s^2), \quad (2.2)$$

where $A_{*\hat{s}}$ denotes a row of the HMM's transition probability matrix and y is the observed input pattern. The dependent state \hat{s} is defined with respect to both the set of parameters for state s and a set of parameters unique to state \hat{s} :

$$A_{*\hat{s}} = (1 - z_{\hat{s}})\vec{a}_s + z_{\hat{s}}\vec{a}_{\hat{s}} \quad (2.3)$$

$$p(y|\hat{s}) \sim (1 - \zeta_{\hat{s}})N(\mu_s, \sigma_s^2) + \zeta_{\hat{s}}N(\mu_{\hat{s}}, \sigma_{\hat{s}}^2). \quad (2.4)$$

The mixing parameters $z_{\hat{s}}$ and $\zeta_{\hat{s}}$ govern the relative contribution of each component. This allows a dependent state to, for example, share the same transition probability vector with its corresponding independent state ($z_{\hat{s}} \approx 0$) but adopt a different distribution of y values ($\zeta_{\hat{s}} \approx 1$). When a mixing parameter is close to zero, the second mixture component plays no role in the likelihood of y under the model, thus reducing the model's effective complexity.

The prior over each mixing parameter was a highly sparse beta distribution that mildly favored smaller values: $z_{\hat{s}} \sim |\text{Beta}(\delta_1, \delta_2)|$ and $\zeta_{\hat{s}} \sim |\text{Beta}(\delta_1, \delta_2)|$. All fixed hyperparameters are listed in Table 1. For the first components' parameters, vague priors were used: $\mu_s \sim N(\xi, \kappa_1^{-1})$, $\sigma_s^{-2} \sim \Gamma(\alpha_1, \beta_1)$, and $\vec{a}_s \sim \text{Dirichlet}(\delta_a, \delta_a, \dots, \delta_a)$, where δ_a specifies the prior over every within-context transition probability. The prior over each $\vec{a}_{\hat{s}}$ was the same as \vec{a}_s . The priors for $\mu_{\hat{s}}$ and $\sigma_{\hat{s}}$ were biased to be somewhat similar to μ_s and σ_s , further reducing the dependent context's contribution to the complexity penalty: $\mu_{\hat{s}} \sim N(\mu_s + h, \kappa_2^{-1})$ and $\sigma_{\hat{s}}^{-2} \sim \Gamma(\alpha_2, \alpha_2\sigma_s^2)$. The purpose of h was to encourage model parameterizations in which $\zeta_{\hat{s}} = 0$ instead of $\zeta_{\hat{s}} = 1$ and $\mu_{\hat{s}} = \mu_{s,1}$. A more sophisticated approach might involve splice sampling from the nonconjugate prior $\mu_{\hat{s}} \sim I(\mu_s, r)N(\mu_s, \kappa_2^{-1})$, where the indicator function I is zero over some interval around the mean,

Table 1: Hyperparameter Values That Define the Prior Distributions over Each Parameter.

HMM Parameters	Prior	Value
Transition probabilities \vec{a}	δ_a	0.8
Parameter sharing probabilities z_s and ζ_s	δ_1, δ_2	0.1, 0.05
Mean of μ_s	ξ	0
Inverse variance of μ_s	κ_1	0.01
Inverse variance of μ_s	κ_2	4
Shape parameter of σ_s^{-2}	α_1	2
Shape parameter of σ_s^{-2}	α_2	10
Rate parameter of σ_s^{-2}	β_1	0.1
Offset of mean of μ_s from μ_s	h	0.4

$[\mu_s - r, \mu_s + r]$; however, the simpler procedure was sufficient for our simulations.

2.4 State Inference and Model Selection. Let M_k denote a model composed of N_S^k states grouped into N_C^k contexts. Let $\theta_k = (\vec{z}, \vec{\zeta}, A, \vec{\mu}, \vec{\sigma}^2)$ denote the parameters of the states in M_k . Given a sequence of noisy inputs, $y_{1,\dots,T}$, the state inference problem is to infer under which sequence of states the inputs are most probable. For state inference, it is assumed that both M_k and θ_k are known and that the HMM is the generative model of $y_{1,\dots,T}$, that is, each HMM state corresponds to an e-state in the world. Thus, inferring the HMM state serves as a proxy for inferring the true e-state. This is traditionally done using the Viterbi algorithm (Viterbi, 1967).

The context learning problem can be understood as part of the more fundamental problem of inferring the generative model of the input patterns: given the sequence $y_{1,\dots,T}$, infer the HMM most likely to have generated it. This can be decomposed into two parts: model parameterization and model selection. Model parameterization involves inferring the most likely set of parameters, θ_k , given an HMM M_k with a known structure. Model selection involves inferring which model structure is most likely out of a set of candidates, typically with different numbers of states.

While model selection in a machine learning context typically involves a single batch analysis, the hippocampus provides some representation of context even in completely novel environments. Multiple models with different numbers of contexts were therefore compared based on initial subsequences, $y_{1,\dots,\tau}$, $\tau < T$. As τ increased, the models were provided with progressively more input patterns to fit, resulting in changes in their likelihoods.

In a Bayesian setting, models are compared based on relative likelihoods, given the observed data $(y_{1,\dots,T})$. Assuming equal prior probability across

models, this can be reduced to calculating the odds of $y_{1,\dots,T}$ under each model:

$$B_{k\ell} = \frac{P(y_{1,\dots,\tau}|M_k)}{P(y_{1,\dots,\tau}|M_\ell)}. \quad (2.5)$$

This is known as the Bayes factor (for review, see Kass & Raftery, 1995), and all model comparisons calculated will be presented as a log Bayes factor. $p(y_{1,\dots,\tau}|M_k)$ is the marginal likelihood of $y_{1,\dots,T}$ under M_k , marginalized over the entire parameter space:

$$p(y_{1,\dots,\tau}|M_k) = \int_{\theta \in \Theta} p(y_{1,\dots,\tau}|\theta_k, M_k) p(\theta_k|M_k) d\theta_k, \quad (2.6)$$

where $p(y_{1,\dots,\tau}|\theta_k, M_k)$ is the conditional probability of the inputs given a particular parameterization of the model, and $p(\theta_k|M_k)$ is the prior probability density over the parameter space.

As mentioned above, equation 2.6 provides an implicit complexity penalty. Importantly, while $p(y_{1,\dots,\tau}|\theta_k, M_k)$ depends on $y_{1,\dots,\tau}$, $p(\theta_k|M_k)$ does not. Since the difference in conditional probability between models typically grows with the length of $y_{1,\dots,\tau}$, the Bayes factor asymptotically determines by how well each model fits $y_{1,\dots,\tau}$. However, for smaller input sequences ($\tau \ll T$), the prior uncertainty over the parameter space, $p(\theta_k|M_k)$, can strongly influence the Bayes factor.

The computational challenge in calculating Bayes' factors is to accurately calculate $p(y_{1,\dots,\tau}|M_k)$, which requires integration over the parameter space (see equation 2.6). A closed-form solution does not exist, but any accurate numerical approximation method is equally acceptable, and the results presented here do not depend on the particular method employed. We use a Monte Carlo integration technique described in the appendix.

3 Simulation Results: Multiple Environments

3.1 Gradual Remapping. Recent data have demonstrated that the development of distinct spatial maps for two environments can be gradual (Tanila, Shapiro, & Eichenbaum, 1997; Jeffery, 2000; Lever et al., 2002). Tanila et al. (1997) found that repeatedly rotating two sets of cues in opposite directions engendered an increase in remapping between the cue configurations over time. Jeffery (2000) found that when the same arena was placed in two different room locations, the number of place fields that differentiated between arena room locations gradually increased over several days. Lever et al. (2002) similarly found that when hunting for food pellets alternately in cylindrical and square arenas, different place cells developed distinct place fields in each arena on different days; two weeks of training were required to achieve complete remapping. In each case, repeated conflicts between cue

configurations were gradually resolved by developing separate contexts for each configuration.

Our modeling of gradual remapping is described with reference to the Lever et al. (2002) study, though it could equally well be interpreted with respect to the others. The arena geometry change is interpreted as a small but consistent difference in the afferent input pattern for which a model with a separate context for each arena will, with enough experience, be better suited.

To model gradual remapping, we constructed input sequences, y_1, \dots, T , generated from alternating sessions in each arena, interposed by time spent on a holding pedestal (during which an experimenter would swap arenas). For computational tractability, two positions within the environment—A and B—were modeled (e.g., one in each half of the arena). The simulated rat spent 2 time steps on the pedestal and 10 time steps in each arena, alternating between the positions in order to remove trajectory differences that might suggest the environments were different. An entire sequence contained 16 environment visits, half in each arena, for a total sequence length of 192 samples.

At each position, the input was perturbed depending on the shape of the arena. For example, at position A, the input had mean $\mu_A - \varepsilon$ in the square arena but $\mu_A + \varepsilon$ in the cylindrical arena. The input standard deviation for all states was $\sigma = 0.125$, and the value of ε used was 0.175, so the input distributions of the two contexts overlapped.

The likelihoods of the one- and two-arena models shown in Figure 1 were compared using the model selection framework described above. The “no pretraining” line in Figure 2 shows that as more experience is gained in each environment, the Bayes factor gradually increases toward a decisive positive value. This suggests that the observed gradual remapping reflects an underlying statistical process: an evidence-based transition to a more complex contextual model.

If one were to pretrain rats by exposing them first to one arena for multiple sessions, how should this pretraining experience affect the rate of remapping? The “ N sessions” lines in Figure 2 reflect a training sequence that starts with N sessions of pretraining in one environment, followed by alternating sessions in the two environments. These simulations predict that pretraining should both delay the onset of gradual remapping and hasten its completion. With few experiences, the second arena looks like a “noisy” version of the first; however, the larger sequence of experiences due to pretraining eventually permits the environments to be distinguished more rapidly.

3.2 Failure to Generalize. In an extension of the Jeffery (2000) study, Hayman, Chakraborty, Anderson, and Jeffery (2003) trained rats in a box placed in two different locations in the room. After place cells gradually learned to remap between box positions, the color of the box and floor was

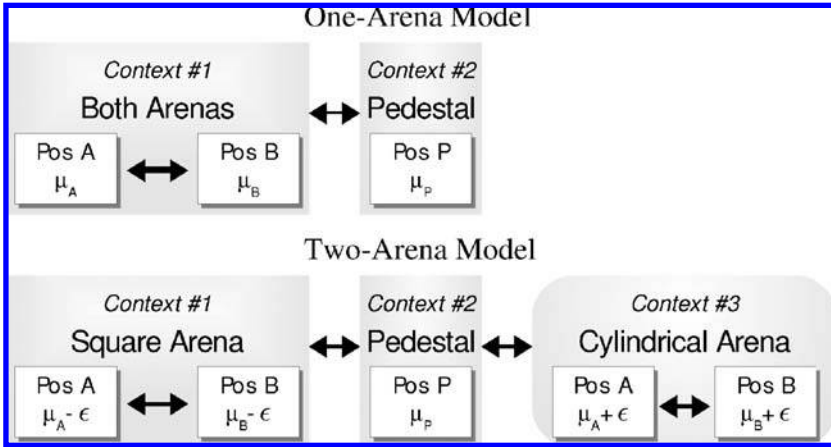


Figure 1: Diagram of one- and two-arena models used to model gradual remapping. The small white boxes represent states in each model, and the large gray boxes indicate how the states are grouped into contexts. Arrows indicate transition probability parameters that are either sampled (within a context) or fixed (between contexts).

changed from white to black. This substantial sensory change resulted in an immediate, complete remapping. Interestingly, despite multiple days of training in a white box in the same two locations, place cells did not immediately discriminate the two locations of the black box.

Their result can be understood as a consequence of the relative differences of the box position and box color manipulations. Specifically, while the position shift created a relatively subtle change, the color change was large. We model their experiment by extending the gradual remapping simulation to include two additional “black box” contexts whose hippocampal inputs are similar to each other but not to states in the original two contexts. Specifically, we constructed input sequences, $y_{1,\dots,T}$, generated from alternating sessions in the white box in room locations 1 and 2, followed by alternating sessions in the black box in room locations 1 and 2. Box visits were, as before, interposed by time spent on a holding pedestal. The two white box positions were modeled identically to the square and cylindrical arenas in the previous simulation. The two black box room locations were modeled in an analogous fashion: the hippocampal inputs for corresponding rat positions in the two black boxes differed by $\pm\epsilon$, but their μ values differed substantially from hippocampal inputs in the white box.

Figure 3 shows the three- and four-context models evaluated on the sequence of hippocampal inputs. Figure 4 shows the evolution of the likelihoods of the three- and four-context models over the course of training in the black box, subsequent to full training in the white box. For comparison,

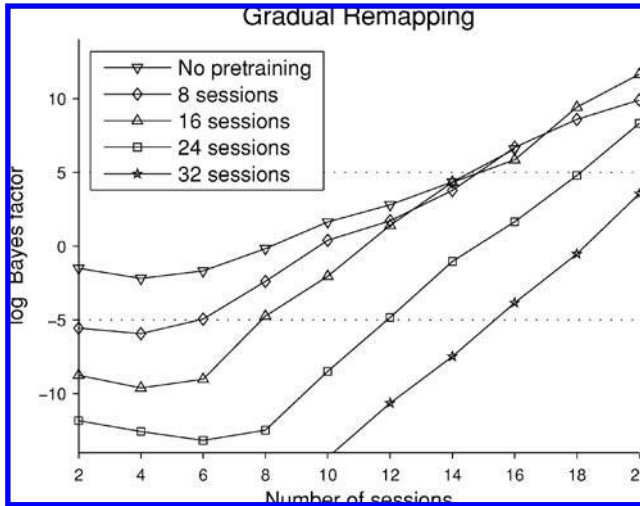


Figure 2: Gradual remapping simulation results, where the log Bayes factor indicates the log likelihood of the two-arena model relative to the one-arena model (see model description in Figure 1). With more experience in each environment, the relative likelihood of the two-arena model increases toward certainty. As a result of first pretraining in one arena, the two-arena model is initially less likely, but preference for this model then increases more rapidly. The dotted lines at ± 5 denote the thresholds beyond which there is essentially no statistical uncertainty in which model is preferred.

training of one- and two-context models in the white box (structured as in Figure 1) is shown as well. Since the white and black boxes differ so strongly, the question of whether the black box locations should be represented as one context or two is statistically unrelated to the representation of the white box locations. The relative likelihood of the four-context model therefore increases at the same gradual rate, predicting that rats should show the same gradual remapping between room locations in both white and black boxes. Consistent with this prediction, Hayman et al. (2003) found in the one rat whose rate of remapping was fastest that the rat remapped between black box locations at the same rate as between white box locations. (The black box sessions were not continued long enough to assess remapping in the rats that were slower to remap in the white box.)

3.3 Morph Environments. J. Leutgeb et al. (2005) trained rats in square and cylindrical environments until the degree of remapping between arenas reached asymptote. Then they exposed the rats to a sequence of arenas whose wall shape systematically morphed between the square and cylinder. They found that the rats' place cell activity in the morph environments reflected various averages of the two maps.

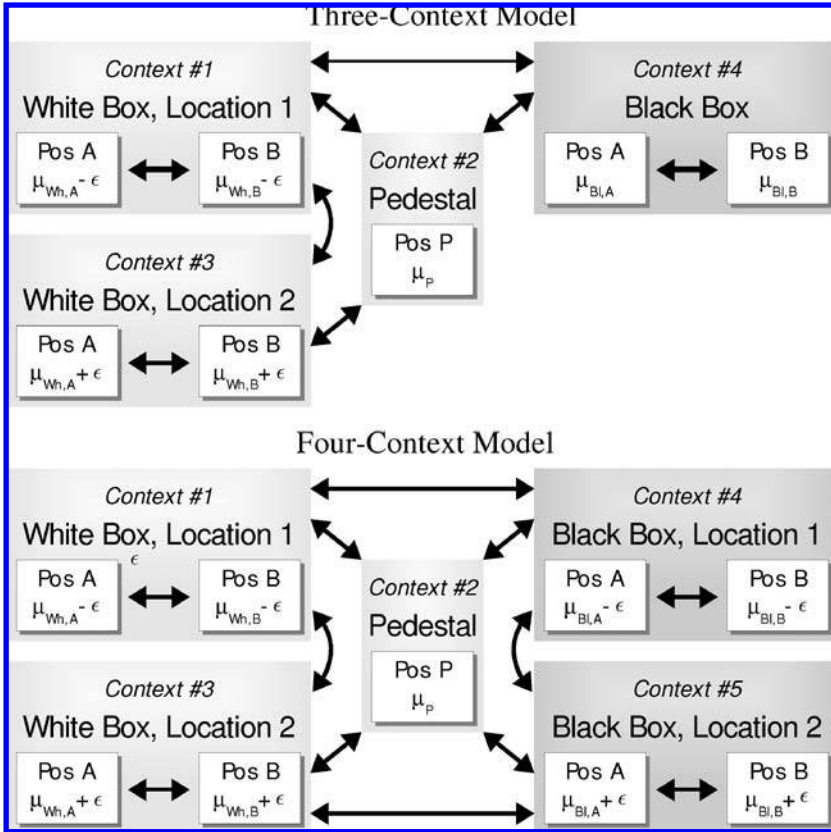


Figure 3: Diagram of three- and four-context models used to model the Hayman et al. (2003) study. The two white box contexts differ subtly, as do the two black box contexts. However, the white box and black box contexts are highly differentiated.

What if the rats were first trained on the morph environments and then on just the square and cylindrical environments? In the gradual remapping simulations, separate contexts for the square and cylindrical arenas develop because they result in distinct clusters of hippocampal input patterns. However, substantial pretraining in the morph environments would generate one broad cluster of input patterns, predicting that morph training should inhibit remapping during subsequent training in just the square and cylindrical arenas.

To model this hypothesized experiment, five arenas were used, and input patterns were generated that varied linearly from the square arena (arena 1) to the cylindrical arena (arena 5). For example, at position *A*, the inputs in the five arenas had means $\mu_A - \epsilon$, $\mu_A - \frac{1}{2}\epsilon$, μ_A , $\mu_A + \frac{1}{2}\epsilon$, and $\mu_A + \epsilon$.

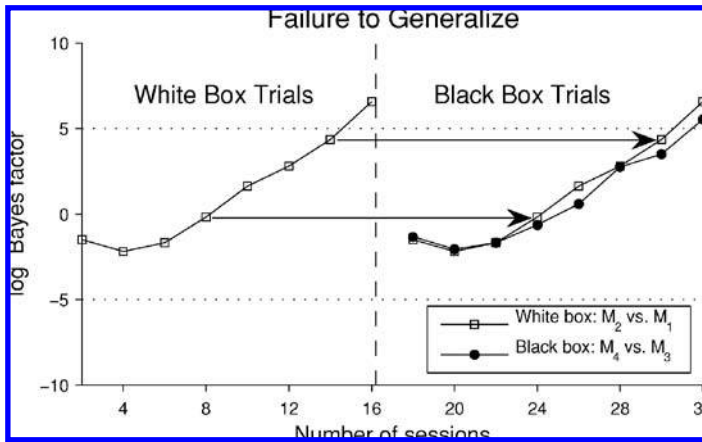


Figure 4: Simulation results for the Hayman et al. (2003) study, where the log Bayes factor indicates the log likelihood of the four-arena model relative to the three-arena model or the two-arena model relative to the one-arena model (see model description in Figure 3). The preference for the four-arena model over the three-arena model (black box) increases at the same rate as the two-arena model over the one-context model (white box). A copy of the white box log Bayes factors is superimposed on top of the black box log Bayes factors to illustrate their similarity.

We constructed input sequences beginning with sessions in just the morph arenas, selecting arena 3 twice as often as arena 2 or 4 to reinforce a single cluster distribution of input patterns. These sessions were followed by alternating sessions in the square and cylindrical arenas. All other details were the same as in the gradual remapping simulation described previously.

Simulation results are shown in Figure 5. Compared to the baseline condition (no morph training), the 16 morph training sessions roughly doubled the time required to complete remapping during subsequent cylinder and square training. With 16 pretraining sessions, remapping was further delayed. This prediction stands in direct contrast to hippocampal models based on independent component analysis (ICA), which would predict place code differentiation during the morph pretraining, as the arena shape acts as an independent source of variation (Lörincz & Buzsáki, 2000).

The disorientation study by Knierim, Kudrimoti, & McNaughton (1995) invites a similar interpretation. Rats were trained to forage for food in a cylinder with an orienting cue card. They found that when rats were disoriented by carrying them around in a closed box before being placed in the arena, their place fields failed to align with the cue card after several sessions, while nondisoriented rats maintained the alignment. They interpreted their results to suggest that the cue card was perceived to be in a different location each time the rats entered the arena after disorientation.




Figure 5: Morph experiment simulation results, where the log Bayes factor indicates the log likelihood of the two-context model relative to the one-context model (see model description in Figure 1). Pretraining in the morph environments inhibits adoption of a two-context model during later training only in square and cylindrical arenas. During morph session training, the Bayes factor decreases, suggesting the morph environments are best represented together as a single context. During subsequent training on square and cylindrical environments, more training sessions are required to justify separate contexts for the two environments than would be required without morph environment pretraining. Simulation results involving morph training were calculated after each block of eight sessions, and simulation results not involving morph training were calculated after each block of four sessions. All plots begin after completion of the first block of training, that is, after four or eight sessions.

After several sessions, the cue card was perceived as “unstable” and therefore was ignored. Moreover, even after multiple subsequent sessions without disorientation, place fields never developed a consistent alignment with the cue card.

Their experiment likely reflects the influence of the head direction system on the place code more than any computational process within the hippocampus itself. Nonetheless, their findings suggest a similar type of statistical inference elsewhere in the brain: if a cue varies in an apparently random manner, its impact on the overall representation of an environment should be minimized.

4 Simulation Results: Reversal Learning

While many conditioning paradigms are not hippocampus dependent, reversal learning has consistently shown dependence on hippocampal

function (Kimble & Kimble, 1965; Silveira & Kimble, 1968; Winocur & Olds, 1978; Berger & Orr, 1983; Neave, Lloyd, Sahgal, & Aggleton, 1994; McDonald, Ko, & Hong, 2002; Ferbinteanu & Shapiro, 2003). Consistent with this interpretation are place cell studies showing context-specific firing patterns during spatial reversal learning tasks (Ferbinteanu & Shapiro, 2003; Smith & Mizumori, 2006).

One of the most interesting aspects of reversal learning is that repeated reward reversals lead to progressively faster behavioral reversal. Two studies have demonstrated that after repeated reward reversals, rats are capable of reversing behavior after a single error trial (Buytendijk, 1930; Duffort et al., 1954). Additionally, Pubols (1962) showed near-perfect reversal performance (fewer than 0.5 errors on average after initial error trial), and Brunswick (1939) showed that even when single-trial reversal performance is not yet achieved, most of the improvement is observed on the second trial. Thus, rats can be trained to select a different (previously learned) behavioral strategy after a single error trial.

Several studies have also explored the impact of partial reinforcement on reversal learning, considering cases in which the "correct choice" is rewarded on only some percentage of trials, as well as cases in which the "incorrect choice" is also rewarded on some smaller percentage of trials (Brunswick, 1939; Wike, 1953; Grosslight, Hall, & Scott, 1954; Elam & Tyler, 1958; Wise, 1962; Pennes & Ison, 1967). The pattern of data across these studies suggests that the more similar the original and reversed contingencies, the more slowly the animal learns the reversal. While intuitive, this suggests that the impact of a trial on a reward association is weighted by how informative the trial is perceived to be. If the expectation of a particular outcome (reward, no reward) is more uncertain, then observing the outcome provides less information about whether the distribution of outcomes has changed.

Previous approaches to modeling reversal learning have posited that the discrimination is relearned during each reversal. Hasselmo et al. (2002) theorized that the hippocampal facilitation of reversal learning was due to quick unlearning and relearning of the association between choice and reward within the hippocampal representation. Learning in their model is unsupervised (Hebbian), and they suggest that reversal learning deficits due to hippocampal or theta modulation impairment are caused by the inability to separate new learning from past associations. Unfortunately, since this model completely relearns the reward association after each reversal, no savings with repeated reversals is predicted. Also, since the current association is dissociated from the previous association, partial reinforcement would not affect the speed of reversal learning under this model.

Another series of models have proposed that the hippocampus plays a role in stimulus representation (Gluck & Myers, 1993, 1996; Myers, Gluck, & Granger, 1995), performing both stimulus compression and predictive stimulus differentiation. A similar model has been proposed by Schmajuk &

DiCarlo (1991, 1992). Both models propose that the hippocampus is critically involved in learning a conjunctive stimulus layer (hidden layer), though their network topologies differ somewhat. In addition, both models train the hidden layer using variants of the backpropagation learning rule. Schmajuk and DiCarlo propose a more direct, biologically plausible implementation, while Gluck and Myers (1993, 1996) suggest only that a functionally similar computational process occurs *in vivo*. Interestingly, "context" is represented in both models as an external sensory input that is integrated with other inputs in the hidden layer rather than as an internal categorization of the input that is inferred.

Both models demonstrate some savings during repeated retraining (on serial reversals or serial extinctions and renewals). In the Gluck and Myers model, hidden layer discrimination increases gradually, making associations with the output layer simpler to learn. In the Schmajuk and DiCarlo model, the learning rate of the input-to-hidden-layer weights is higher, so the increased weighting of the hidden layer representation at the output layer over the repeated reversals leads to faster relearning. In either case, both models fundamentally rely on some degree of retraining during each reversal. However, the single trial reversals observed experimentally occurred based on an error trial alone (Buytendijk, 1930; Dufort et al., 1954). Even with an arbitrarily high learning rate, it is not clear how one could retrain a backpropagation network without at least one positive trial.

In addition, the partial reinforcement conditions are not well explained by these models, or other delta-rule based models (e.g., Rescorla & Wagner, 1972; Pearce & Hall, 1980), since learning is not adjusted based on the entropy of the trial's outcome. Similarly, none of these models accounts well for the partial reinforcement extinction (PRE) effect, in which extinction training is prolonged following partial reinforcement. With generalized delta rule learning, partial reinforcement results in a weaker association between behavioral choice and outcome. During a reversal (or extinction), this weaker association would be easier to expunge than a stronger association would be. However, the data show the opposite results. A few recent statistical models account well for the basic PRE effect (Gallistel & Gibbon, 2000; Courville, 2006), but they do not address the progressive improvement observed in reversal learning.

An alternative interpretation of the reversal learning data, as proposed by Hirsh (1974), is that the hippocampus represents each reward condition as a separate context. If the reward contingencies of both contexts are represented simultaneously, then after initial training, no retraining should be required after each reversal. Rather, recent trial outcomes can be compared to prior knowledge of each context to infer which context is active. Similarly, "sequential theory" posits that the pattern of prior rewards forms a context that influences which conditioned associations are retrieved from memory (Capaldi, 1994; Capaldi & Neath, 1995).

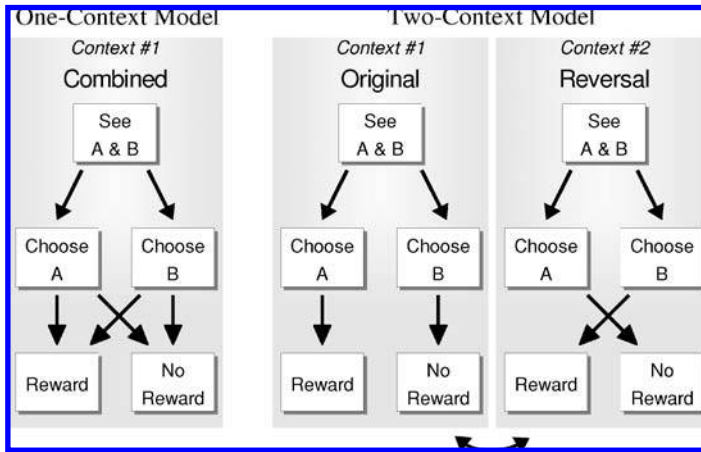


Figure 6: Diagram of one- and two-context models used to model reversal learning. Arrows within each context indicate transition probabilities that are typically high; however, transitions between any two states are possible.

In a full-reinforcement reversal learning paradigm, knowledge of both contexts allows a single error trial to be sufficient to infer a context switch, since the choice during the error trial should yield no reward only in “the other” context. One might think of this process as analogous to the problem of self-localization in spatial navigation, a function that has also been attributed to the hippocampus (Touretzky & Redish, 1996). The statistical formulation of context learning also correctly quantifies the increased uncertainty in the adoption of a new context under partial reinforcement conditions.

Interestingly, if the reversal is performed in a different environment, the environmental cues substantially improve adaptation to the reversal, even in hippocampal animals (McDonald et al., 2002). This finding reinforces the notion that what the hippocampus contributes to reversal learning is a contextual cue to separate the two discriminations.

To model reversal learning, we constructed a simple model of discrimination learning (see Figure 6). The discrimination begins at the start state, denoting the availability of two options (A and B). The start state leads to either of two states reflecting the rat’s choice. The choice states in turn lead to reward and no reward states, indicating the outcome of that choice. (Such reward states are justified by studies showing reward-related responses in the hippocampus; Tabuchi, Mulder, & Wiener, 2003; Smith & Mizumori, 2006.)

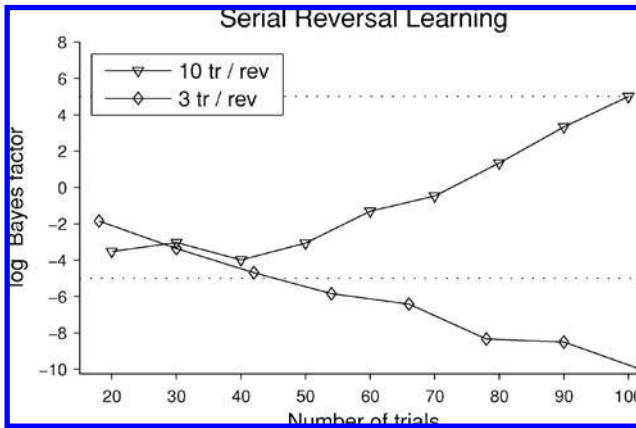


Figure 7: Serial reversal learning simulation results, where the log Bayes factor indicates the log likelihood of the two-context model relative to the one-context model (see model description in Figure 6). With only three trials per reversal, the Bayes factor decreases, suggesting that rapid reversals are indistinguishable from partial reinforcement. Less frequent reversals (10 trials per reversal) lead to an increase, supporting the adoption of a separate context for each reward contingency.

4.1 Serial Reversal Learning. To simulate serial reversal learning training, alternating blocks of 10 trials were generated in which, for odd blocks, choice A was rewarded, and, for even blocks, choice B was rewarded. To simulate behavioral learning, the rat's choice was selected randomly such that during the first block, the probability of selecting the correct choice exponentially approached 58.5% from 50%; during the second block, the correct choice probability exponentially approached 62% from 41.5% ($100\% - 58.5\%$); during the third block, the correct choice probability exponentially approached 65.5% from 38%. This continued for 10 blocks, increasing the final correct choice probability in each successive block by 3.5%. The exponential "decay" of choice probabilities within each trial approximates the trial-by-trial error data of Brunswick (1939).

One- and two-context models, shown in Figure 6, were compared as the number of training blocks was increased; the results are shown in Figure 7. With increased training, the likelihood of the two-context model gradually increases, predicting a gradual adoption of context-specific hippocampal firing patterns. Smith and Mizumori (2006) showed that context specificity developed during reversal learning training, though they do not examine the time course of remapping across training sessions. Figure 8 illustrates context selection in the two-context model, demonstrating how a context switch can be inferred after a reversal.

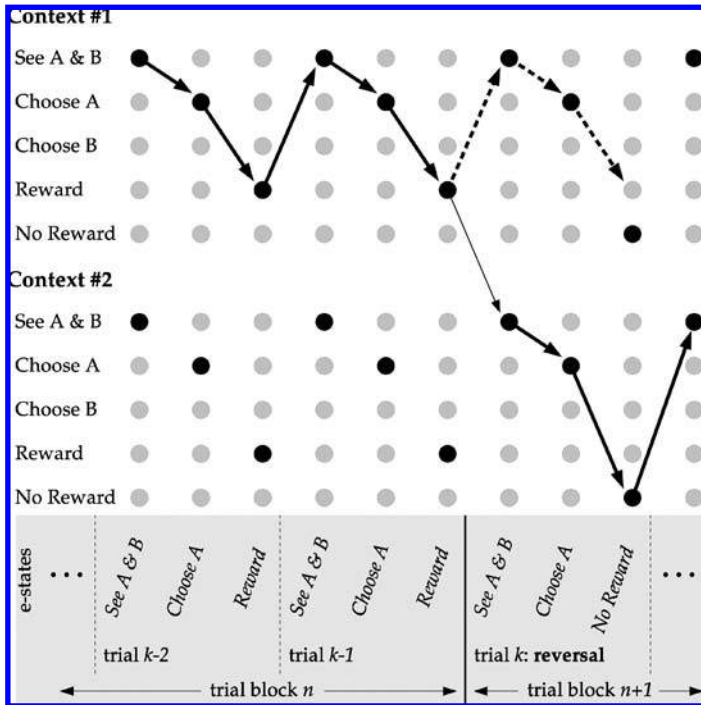


Figure 8: Context selection in the two-context model. Each circle represents a state in the HMM, and each column represents the full set of states of the two-context model at a particular point in time. The gray box shows e-states (hippocampal inputs) from three trials. Trials $k - 2$ and $k - 1$ are at the end of a block of trials in which choice A is rewarded. Trial k is the first trial of a new block in which choice B is rewarded. HMM states that are probable given the hippocampal input are darkened. With full reinforcement, only one choice in each context is associated with reward. Therefore, while the dashed-line path (remaining in Context #1) initially appears to be more likely, the lack of reward indicates that trial k is actually part of Context #2. Once the context switch has been inferred, the most likely path through the states will continue within Context #2 until choice B is no longer rewarded.

If the number of trials per block is substantially reduced, the reward structure becomes difficult to distinguish from partial reinforcement. When training sequences were generated comprising three trials per block, the two-context model required too many context switches to be justified, since context switches are constrained to be unlikely. This resulted in the progressively decreasing likelihood of the two-context model, as shown in Figure 7.

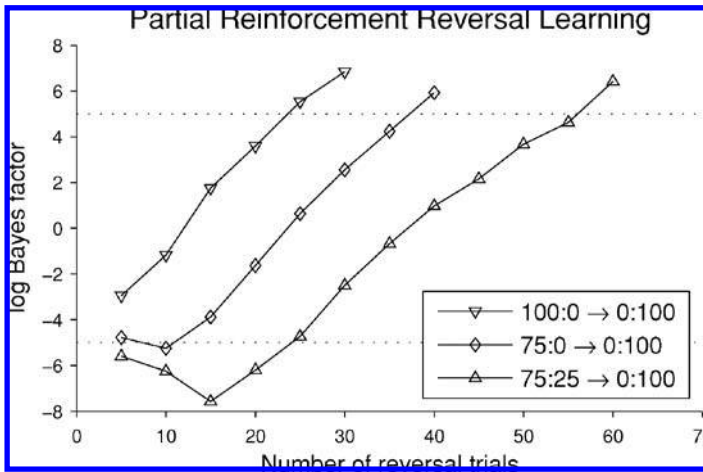


Figure 9: Partial reinforcement reversal learning simulation results, where the Bayes factor indicates the log likelihood of the two-context model relative to the one-context model (see model description in Figure 6). When the similarity of the two reward contingencies is increased, more trials are required to justify representing them as separate contexts.

4.2 Partial Reinforcement and Reversal. In order not to conflate the effects of partial reinforcement during the original and reversal discriminations, the training paradigm typically applies partial reinforcement to choices during the original discrimination, leaving the reversal condition unambiguous (Wike, 1953; Grosslight et al., 1954; Elam & Tyler, 1958; Wise, 1962; Pennes & Ison, 1967). We consider three cases: full reinforcement (100:0), in which choice A is always rewarded and choice B is never rewarded; partial reinforcement of A (75:0), in which choice A is rewarded only 75% of the time; and partial reinforcement of A and B (75:25), in which choice B is also rewarded 25% of the time. In all cases, only choice B is rewarded during the reversal discrimination (0:100).

Training consisted of a single original and single reversal block, the length of each being equal and set so that the Bayes factor would be roughly 5.0 by the end of training on both blocks. The rat's choice was selected probabilistically, where the probability of a correct choice exponentially approached 90% from 50% in the original block and 90% from 10% in the reversal block. This simulated behavior was designed to mimic the progressive bias animals show in favor of the rewarded choice over time (e.g., Wike, 1953).

Figure 9 shows that as the original discrimination is made more ambiguous, learning to differentiate it as a separate context from the reversal discrimination requires progressively more trials. We therefore predict that

context-specific hippocampal representations would develop more slowly in these partial reinforcement paradigms.

5 Simulation Results: Sequence Learning

The hippocampus has been implicated in a variety of tasks involving sequences (Kesner & Novak, 1982; Chiba, Kesner, & Reynolds, 1994; Gilbert, Kesner, & Lee, 2001; Agster, Fortin, & Eichenbaum, 2002; Fortin, Agster, & Eichenbaum, 2002; Kesner, Gilbert, & Barua, 2002). For example, Fortin et al. (2002) found that after being presented with a random sequence of odors, hippocampal rats could not choose the earlier of two odors from the sequence. Of particular interest has been the study of overlapping sequences, in which multiple sequences share common middle elements but distinct beginning and ending elements. Agster et al. (2002) found that while hippocampal rats could disambiguate two partially overlapping olfactory sequences, intertrial interference or a delay condition could impair their performance substantially.

Several models have proposed that the hippocampus develops separate contextual representations of each sequence that serve to associate the ambiguous middle elements with the rest of the sequence (Levy, 1989, 1996; Wallenstein & Hasselmo, 1998). These models predict that if a rat were to repeatedly travel down a common maze arm that was part of two different paths (e.g., a continuous figure 8 pattern), place cells on the common arm would fire differently depending on which path the animal was traveling.

Some studies have confirmed this finding (Frank, Brown, & Wilson, 2000; Wood et al., 2000), while others found no path-related differences (Lenck-Santini et al., 2001; Hölscher et al., 2004). Intriguingly, Bower et al. (2005) were able to reproduce both cases by varying the shaping procedure used to train the rats. The studies that found no path-related differences nonetheless reported that rats were able to learn the task, a result consistent with Ainge and Wood (2003), who found that hippocampal lesions did not impair the continuous version of the task. However, when a small delay was added at the start of the common maze arm, Ainge and Wood found that hippocampal rats were impaired. Subsequently, two groups trained unlesioned rats on a figure 8 maze, each finding, paradoxically, that path-specific modulation of place fields on the common maze arm occurred with no delay (when the task was not hippocampus dependent) but largely disappeared when a delay was added (Ainge, van der Meer, & Wood, 2005; Robitsek, Fortin, & Eichenbaum, 2005). Despite this disappearance, rats could still perform the task. As Bower et al. (2005) point out, the sequence dependence of place cell activity is likely attributable to differing input from some extrahippocampal brain area instead of a contextual representation developed within the hippocampus.

An alternative hypothesis, adopted by Hasselmo and Eichenbaum (2005), is that sequence replay (Foster and Wilson, 2006) in the hippocampus

guides behavior independent of whether path-specific remapping is seen on the common arm. Specifically, an extra-hippocampal brain area incrementally learns rules for completing each sequence, given its beginning; replay of the beginning of the current sequence is then sufficient to complete the task. (With a higher learning rate, Hasselmo et al., 2002, would likely provide an elegant model of such one-shot learning.) Such a division of labor explains why Agster et al. (2002) found evidence of rats learning overlapping sequences but failing, under some circumstances, to recall which sequence had most recently begun.

The rapid alternation between sequences on a figure 8 maze is incompatible with our constraint that context switches should be rare. To demonstrate this, we constructed a simulation of the figure 8 task. Five positions on the track were modeled and were traversed in a six-step loop: start left, center, end right, start right, center, end left, repeat. A trial constituted a pass through one start arm, the center arm, and the opposing end arm. One- and two-context models (see Figure 10) were compared, and the results are shown in Figure 11. Even after just 12 trials, the two-context model is astronomically unlikely.

To demonstrate that the context switch penalty in the two-context model is the specific cause of the low Bayes factors, the one-context model was also compared with the generative model (see Figure 10). Figure 11 shows that decisive preference for the generative model is attained by 30 trials.

Bower et al. (2005) have considered training regimes that promote sequence disambiguation of the common path, presumably due to afferent input from another brain region. However, were rats to be trained repeatedly on one sequence and then the other, our framework would predict that sequence-specific encodings of the common path would reliably develop due to intra-hippocampal mechanisms sensitive to temporal mismatch.

6 Discussion

Advances in Bayesian computational statistics techniques over the past decade have opened the door to evaluating the marginal model probabilities of many new classes of models. As researchers use these techniques, insights into animal learning and human reasoning have begun to arise from their formulation as Bayesian inference problems, often over multiple generative models that provide competing explanations of a corpus of data (Tenenbaum & Griffiths, 2001; Courville, Daw, Gordon, & Touretzky, 2003; Courville, Daw, & Touretzky, 2004; Griffiths & Tenenbaum, 2005; Daw, Niv, & Dayan, 2005). In a similar vein, work presented here advances a theory of context learning in which hippocampal input patterns are grouped together in the same context when they reliably cluster together in time. Choosing between generative models that group the experiences into different numbers of clusters is therefore the fundamental challenge of context learning.

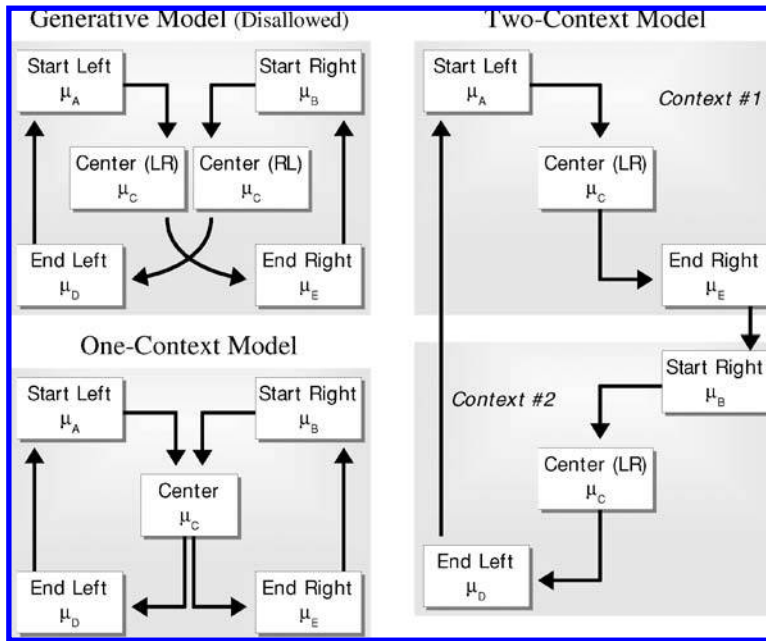


Figure 10: The generative model and one- and two-context models used to model sequence learning. Note that the generative model includes two states (Center LR and Center RL) in the same context that represent the same distribution of input patterns (identical μ_C values). Our weak-identifiability constraint (see section 2.1) that such states can be represented only when in separate contexts excludes the generative model from the class of possible hippocampal context models. Arrows within each context indicate transition probabilities that are typically high; however, transitions between any two states are possible.

Sequence learning models (Levy, 1989, 1996; Wallenstein & Hasselmo, 1998) provide a rather different notion of context, one oriented toward binding together elements of a temporal sequence. Interestingly, this binding process simultaneously serves to differentiate elements based on their order within the sequence. By contrast, in our theory, context learning groups sequence elements together without disambiguating multiple occurrences of the same element. Thus, sequence learning models predict that alternating, overlapping sequences should be represented with different contextual bindings, whereas our theory groups them together into one context. The failure for sequence-dependent hippocampal representations to be consistently observed (e.g., Bower et al., 2005) or to serve a behavioral function (Ainge & Wood, 2003) argues against such representations playing a critical role in the disambiguation of alternating, overlapping sequences. Others have argued that such sequential encodings are instead formed in medial




Figure 11: Sequence learning simulation results, where the log Bayes factor indicates log likelihood of the two-context model relative to the one-context model (see model description in Figure 10). The relative likelihood of the two-context model decreases precipitously with experience due to the constraint that transitions between contexts are of low probability. If the generative model were admissible, it would be adopted over the one-context model, as its log likelihood relative to the one-context model increases to certainty with experience.

temporal lobe areas outside the hippocampus proper such as the entorhinal cortex (Howard, Fotedar, Datey, & Hasselmo, 2005).

An important aspect of this theory is that it distinguishes between the inference processes of context learning and context selection: context learning may be gradual over many days, while context selection should be an abrupt process. Multibasin attractor models (Samsonovich & McNaughton, 1997; Redish & Touretzky, 1998; Tsodyks, 1999; Dobioli et al., 2000) have demonstrated how multiple contexts, each a stable basin of attraction, could be simultaneously represented within the same network. In such models, context selection involves restabilization of the network in the most appropriate basin. What these models lack is an explanation of the gradual development of new contextual representations.

By contrast, backpropagation models (Schmajuk & DiCarlo, 1992; Gluck & Myers, 1993) have attempted to address the gradual development of new (hidden layer) representations. However, these networks do not have multistable activity patterns; they have no “memory” of the current context across time beyond what is encoded in the weights. Thus, they fail to capture the one-trial context switching behavior observed in reversal learning. In addition, backpropagation and similar delta rule learning models do not properly adapt learning to the information provided by each trial and

therefore cannot account for the slower reversal of partial reinforcement reward contingencies.

Finally, while backpropagation is an efficient search technique for learning high-dimensional mappings, the complexity of the function represented by the neural network is not explicitly considered. (In a machine learning setting, overcomplexity can result in “overfitting,” which is typically detected by cross-validation of the model on a separate data set.) Similarly, independent component analysis (used in a hippocampal model by Lörincz & Buzsáki, 2000) does not adjust the number of independent components based on the observed data. In contrast, the framework here considers the inherent trade-off between model fit and model complexity that underpins any model selection problem. Our framework can most clearly be dissociated from other models that do not consider this trade-off by determining whether remapping is observed subsequent to training in a set of morph boxes: our simulations predict that remapping should not occur.

6.1 Localizing Contextual Representations Within the Hippocampus.

Since our theory concerns contextual representations in the hippocampus, we discuss how it relates to known facts about hippocampal anatomy and physiology. First, we argue that abrupt and gradual remapping are mediated by distinct physiological processes. Specifically, whereas abrupt remapping is caused by a change in the path integrator representation located in dorsal medial entorhinal cortex (dMEC), gradual remapping is caused by experience-dependent representational changes within the DG/CA3 network. Then we discuss the evidence for pattern separation and completion in the DG/CA3 network and how such mechanisms could underpin context learning. We contrast the role of the DG/CA3 network with the role of CA1 in gating the projection of the DG/CA3 representation to efferent cortical areas.

6.1.1 Abrupt Remapping and Attractor Dynamics. Marr (1971) first proposed that the architecture of the hippocampus is well suited to encode new memory traces using orthogonalized representations, which minimize interference between stored patterns. Recordings from hippocampal pyramidal cells confirmed that sparse, orthogonalized representations are formed to encode different places and other features within an environment (O’Keefe & Dostrovsky, 1971; Wood, Dudchenko, & Eichenbaum, 1999), as well as different environments as a whole (Muller & Kubie, 1987). Attractor models of hippocampal function (Samsonovich & McNaughton, 1997; Redish & Touretzky, 1998; Doboli et al., 2000) which address the remapping of place fields between environments, were born out of the observation that, when place cells remap, they appear to remap together. For example, Bostock, Muller, & Kubie (1991) found that the change of a cue card’s color sometimes caused a remapping, and, when observed, all simultaneously recorded cells remapped. When introduced repeatedly to the same

environment, rats with deficient LTP sometimes recalled a completely different map (Barnes, Suster, Shen, & McNaughton, 1997; Kentros et al., 1998).

One difficulty in interpreting such remapping data is in disambiguating the role of the path integrator, currently believed to be in dorsal medial entorhinal cortex (dMEC) (Fyhn, Molden, Witter, Moser, & Moser, 2004; Hafting, Fyhn, Molden, Moser, & Moser, 2005; Fuhs & Touretzky, 2006; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006), from the place code. As noted by Touretzky and Redish (1996), failure to reset the path integrator should cause a substantial change in the afferent input to the hippocampus, resulting in abrupt hippocampal remapping independent of any attractor dynamics in the hippocampus. This reset failure likely underpins the results of S. Leutgeb et al. (2005), who showed that switching between two rooms causes a different and more radical form of remapping than switching between arena shapes in the same room. While dMEC grid cells change phase when the arena is placed in a novel room (Hafting et al., 2005), preliminary findings by J. Leutgeb et al. (2006) show that grid cell phases remain constant when the arena changes shape in the same room, while DG and CA3 undergo rate remapping (see also Quirk, Muller, Kubie, & Ranck, 1992). Thus, abrupt remapping, even when delayed from the first exposure to the environment (e.g., Bostock et al., 1991; Brown & Skaggs, 2002), is likely attributable to a reset failure of the path integrator, which causes remapping simultaneously in all subfields of the hippocampus. By this interpretation, delayed, abrupt remapping reflects a stochastic process where, for rats that attend to the environmental change, there is a fixed probability of PI reset failure on each visit to the perturbed environment, which should result in an exponentially distributed time to first remapping.

6.1.2 Gradual Remapping, Context Learning and the DG/CA3 Network. In contrast to abrupt remapping, the gradual differentiation of contextual representations should be attributable to circuitry within the hippocampus, specifically DG and CA3. Several authors have suggested that exposure to a novel context results in an orthogonalized representation being formed in DG, which is then propagated to CA3 (Marr, 1971; McNaughton & Morris, 1987; Treves & Rolls, 1992, 1994; O'Reilly & McClelland, 1994). Preliminary evidence suggests that the rate remapping between similar contexts observed by S. Leutgeb et al. (2005) originates in DG (J. Leutgeb et al., 2006). In a familiar context, perforant path input and recurrent collaterals in CA3 guide the recall of the previously learned contextual representation. O'Reilly and McClelland (1994) explored a simplified model of DG and CA3, showing that although similar patterns could be mapped to an even more similar CA3 representation (pattern completion), more strongly differing input patterns could be separated as well (pattern separation).

The neural mechanisms of pattern separation and completion resemble the statistical process of clustering: map noisy input patterns into more

similar representations to denote their association with the same cluster; map input patterns associated with different clusters into more distinct representations. One might therefore think of a clustering neural network as an extension of the O'Reilly and McClelland (1994) model in which the threshold between separation and completion is not static, but dynamically adjusted based on the distribution of input patterns. While O'Reilly and McClelland only elucidated the benefits of CA3 perforant path plasticity, NMDA-dependent synaptic plasticity is well known to exist within DG and the CA3 recurrent collaterals. In addition, there is intriguing new evidence that mossy fiber synapses show heterosynaptic plasticity, though changes in synaptic efficacy appear to depend on neighboring synapses in stratum radiatum instead of the depotentiation of the postsynaptic cell (Schmitz, Mellor, Breustedt, & Nicoll, 2003). We propose that one function of this learning is to adjust the separation-completion threshold, gating when new contextual representations would be propagated from DG to CA3.

In a familiar context, once the path integrator and any other brain state is reset, the activity patterns projected from DG onto CA3 should match the patterns projected from the perforant path and recurrent collaterals, modulo some noise. However, if this familiar context is perturbed into a similar but distinct second context, then there should be some mismatch between the pathways: the representation in DG should more accurately reflect the current input patterns, whereas the perforant path and recurrent circuitry should reinforce a previously learned representation. Critically, our framework suggests that if the differences between recalled and input patterns are small, the impact of the DG representation should be minimal, and, if such differences do not repeat, the impact of DG should remain minimal. However, if the same differences are observed repeatedly, the impact of DG on the CA3 representation should increase, causing remapping in CA3 (Fuhs & Touretzky, 2000). In this way, more complex context models may be adopted.

Small but repeated differences would be expected to cause incremental changes in synaptic connectivity in the DG perforant path and mossy fiber pathway to strengthen the impact of DG on CA3. Interestingly, perforant path plasticity in DG can last for months (Abraham, Logan, Greenwood, & Dragunow, 2002), suggesting that this pathway is capable of accruing changes in synaptic efficacy over many training days. However, granule cell neurogenesis causes cells to be gradually replaced, slowly fading away previous associations (Feng et al., 2001). These opposing forces provide a natural balance for input pattern density estimation. New input patterns can be registered, and their impact can be strengthened with repeated exposure; however, this strengthening is tempered by the replacement of granule cells to clear out old memories. It follows from this proposal that DG principal cells should fire at a lower rate in a novel environment and, with repeated exposure, gradually increase their rates as the environment becomes familiar.

While presenting a neural network model of context learning is beyond the present scope, recent physiology, gene expression, and lesion studies are consistent with the proposal that a neural instantiation of context learning should be localized to the DG/CA3 network, whereas CA1 integrates the experience-dependent DG/CA3 representation with the entorhinal cortical representation. Both pattern separation and pattern completion have been observed in CA3 in response to changes in environmental cues (Vazdarjanova & Guzowski, 2004; Lee, Yoganarasimha, Rao, & Knierim, 2004; Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004; Guzowski, Knierim, & Moser, 2004). Vazdarjanova and Guzowski (2004) found in an immediately-early gene expression study that similar environments were represented with a greater similarity in CA3 than CA1, while two very different environments were represented with less similarity in CA3 than CA1. Both Lee et al. (2004) and S. Leutgeb et al. (2004) found that the CA1 representation more directly reflected the current sensory environment, whereas the CA3 representation reflected either pattern completion of a previously learned contextual representation (Lee et al., 2004) or pattern separation to create a new contextual representation (S. Leutgeb et al., 2004).

A behavioral role of CA3 pattern completion is suggested by Nakazawa et al. (2002), who found that performance in a cue-degraded version of the Morris water maze is impaired by CA3 NMDA receptor knockout. Lee and Kesner (2002, 2003) showed that delayed nonmatch to place (DNMP) was impaired in a novel (but not familiar) environment by CA3 NMDA inactivation or DG or CA3 lesion. These deficits may be interpreted as a failure to recall or learn a conjunctive representation of position and target (e.g., hidden platform, object) that could be retrieved via pattern completion using the target as an autoassociative memory cue. (Evidence for such a target representation has been found in prelimbic/infralimbic cortex; Hok, Save, Lenck-Santini, & Poucet, 2005.) Consistent with this hypothesis, DG appears necessary to create such conjunctive representations: DG lesions reduce performance to chance on a working-memory water maze task in which the platform is moved to a new location each day (Xavier, Oliveira-Filho, & Santos, 1999). In relation to our theory, these data point to DG and CA3 to construct a model-based representation of the animal's experiences, including various conjunctive associations instrumental in solving behavioral tasks.

CA1 appears to relay a composition of the model-based CA3 representation and the context-free entorhinal cortical information to efferent cortical areas. Hasselmo and colleagues have presented a series of models and pharmacological data supporting the notion that increased cholinergic modulation decreases the contribution of CA3 to the CA1 representation, but increases plasticity of both the CA3 recurrent and Schaffer collaterals (Hasselmo & Schnell, 1994; Hasselmo, Schnell, & Barkai, 1995; Hasselmo, Wyble, & Wallenstein, 1996). More recently, Yu and Dayan (2005) have proposed a theory of acetylcholine and noradrenaline in which acetylcholine

represents expected uncertainty, whether due to the context being new or to known unpredictability within a familiar context.

Taken together, these models suggest that when the context is informative (low ACh), CA1 should be influenced by CA3; when the context is less informative (high ACh), that influence should be reduced. This gating of the CA3 representation has been confirmed experimentally in novel environments: CA1 shows a stable representation while the CA3 representation evolves over the course of 20 to 30 minutes (S. Leutgeb et al., 2004). Additionally, several monoaminergic neurotransmitters have been implicated in modulating the balance of CA3 and EC input to CA1 (Otmakhova & Lisman, 1998; Otmakhova, Lewey, Asrican, & Lisman, 2005).

The complementary roles of the DG/CA3 and CA1 networks provide some insight with which to interpret discrepancies between the double rotation experiments of Shapiro et al. (1997) and Lee et al. (2004). When local and distal cues were rotated in opposite directions, both studies observed "heterogeneous" or "discordant" responses. However, Shapiro et al., recording mostly from CA1 at the beginning of the experiment, initially observed many more place fields to rotate with the distal cues than the local cues. In fact, the ratio of place fields rotating with each set of cues observed by Shapiro et al. much more closely resembles the ratio of place fields in CA3 tied to each set of cues observed by Lee et al. (2004), suggesting that in the Shapiro et al. study, CA3 strongly influenced the representation in CA1. Shapiro et al. repeatedly trained rats on two standard condition sessions and a single double rotation session (local and distal cues rotated 180 degrees apart) each day, as well as various other less frequent cue scrambling and deletion probe trials. They observed that over time, cells (predominantly in CA3 by the end of the experiment) more strongly remapped between standard and double rotation conditions. This change in the degree of remapping could reflect either the change in cell populations they recorded from or an experience-dependent effect; they do not address this issue statistically. Nonetheless, if we assume that the initial primacy of distal influence on CA1 place cells reflects pattern completion in CA3, then the increase in remapping reflects an experience-dependent transition in CA3 between pattern completion and pattern separation in order to adopt separate contexts for the standard and double rotation cue configurations. This is consistent with our simulations of gradual remapping, which show that two repeatedly experienced and distinct conditions should be differentiated by context. By contrast, Lee et al. trained rats equally on four rotation angles in addition to the standard condition, a training paradigm more akin to our morph experiment simulations. Though the relatively short duration of training by Lee et al. prevents any decisive conclusions, their observation that CA3 maintained pattern completion throughout the course of their experiment (I. Lee & J. J. Knierim, personal communication to M. Fuhs, June 2006) is consistent with our morph experiment simulation results that predict that experience-dependent remapping should not occur in such a case.

6.2 Future Work. While multi-unit recordings from CA3 and CA1 have provided great insight into their differing function, data from DG and the hilus are only beginning to become available. As the activity patterns of cells in these areas are explored, so too does the exploration of physiologically based models of context learning become more tenable.

The framework presented here assumes a perfect memory of past experiences, an assumption certain to be untrue. It would be interesting to explore to what extent memory limitations, perhaps imposed by the representation of the full history of experiences by a set of (in)sufficient statistics, would affect the predictions made by the present framework. Also, the framework here models context transitions as a Markov process with fixed transition probability, which results in the assumption that context transitions occur according to a Poisson distribution. However, Daw, Courville, and Touretzky (2006) have suggested that a more sophisticated semi-Markov process better explains physiology studies of the dopamine system. Context learning *in vivo* may similarly involve a more explicit inference about the amount of time an animal expects to remain in each context.

Appendix: Estimating Marginal Model Probabilities

For these simulations, equation 2.6 was numerically approximated by Monte Carlo integration of the parameter samples using an importance distribution. The importance distribution was constructed in an automated fashion using 500 samples from the posterior parameter distribution, $p(\theta_k | y_{1,\dots,\tau}, M_k)$, and the Gibbs kernels from the Markov chain Monte Carlo (MCMC) sampler that generated the parameter samples (for details, see Frühwirth-Schnatter, 2004). Monte Carlo integration becomes unstable when the tails of the importance distribution are narrower than the posterior distribution along any dimension of the parameter space. To guard against this, the variances of the μ and σ^{-2} components were doubled, and the variance of the a_s component was increased by multiplying each Dirichlet parameter by 0.75.

The MCMC sampler was constructed based on previous sampling techniques for standard HMMs (Chib, 1996). Briefly, the Markov chain was constructed from Gibbs kernels that sample sequentially from both the parameters of the model and latent indicator vectors, which indicate, for each mixture distribution within the model and each time step t , which mixture component is implicated in the observed value y_t . For the shared parameter HMMs, there are three indicator vectors. The vector $S_{1,\dots,T}$ indicates the active HMM state at each time step and was sampled using a standard HMM Gibbs move:

$$p(S_t = i | \dots) \propto A_{i,S_{t-1}} p(y_t | \mu_i, \sigma_i^2) A_{S_{t+1},i}. \quad (\text{A.1})$$

The vector $u_{1,\dots,T-1}$ indicates which component of the transition probability mixture (see equation 2.3) for the current state contributed to each transition. The vector $v_{1,\dots,T}$ indicates which component of the hippocampal input value mixture (see equation 2.4) for the current state contributed to each observation. Gibbs moves were as for standard mixtures:

$$p(u_t = \text{Dep} | \dots) \propto A_{S_{t+1}, \hat{s}} / (A_{S_{t+1}, s} + A_{S_{t+1}, \hat{s}}) \quad (\text{A.2})$$

$$p(v_t = \text{Dep} | \dots) \propto p(y_t | \mu_{\hat{s}}, \sigma_{\hat{s}}^2) / (p(y_t | \mu_s, \sigma_s^2) + p(y_t | \mu_{\hat{s}}, \sigma_{\hat{s}}^2)). \quad (\text{A.3})$$

The values of u_t and v_t are defined only for times when the HMM state indicator vector indicates a state in a dependent context.

Gibbs moves for the HMM parameters were complicated by sharing of parameters between independent and dependent states. Gibbs moves for the mixing parameters were

$$p(z_{\hat{s}} | \dots) \propto \text{Beta}(\delta_1 + n_s^u, \delta_2 + n_{\hat{s}}^u) \quad (\text{A.4})$$

$$p(\zeta_{\hat{s}} | \dots) \propto \text{Beta}(\delta_1 + n_s^v, \delta_2 + n_{\hat{s}}^v), \quad (\text{A.5})$$

where $n_s^u = \#(u_t^{\hat{s}} = \text{Ind})$, $n_{\hat{s}}^u = \#(u_t^{\hat{s}} = \text{Dep})$, $n_s^v = \#(v_t^{\hat{s}} = \text{Ind})$, and $n_{\hat{s}}^v = \#(v_t^{\hat{s}} = \text{Dep})$. The counting function $\#()$ returns the number of occurrences over a time-indexed vector in which the specified condition is satisfied. Gibbs moves for the transition probabilities were

$$p(\bar{a}_{s_1} | \dots) \propto \text{Dirichlet}(\delta + n_1^{\text{ind}}, \delta + n_2^{\text{ind}}, \dots) \quad (\text{A.6})$$

$$p(\bar{a}_{\hat{s}_1} | \dots) \propto \text{Dirichlet}(\delta + n_1^{\text{dep}}, \delta + n_2^{\text{dep}}, \dots). \quad (\text{A.7})$$

The transition probabilities for a state in an independent context, s_1 , were updated based on a transition count, $n_{s_2}^{\text{ind}}$, which counts the number of occurrences of a transition from s_1 to s_2 . If the independent context had a dependent context (with states \hat{s}_1 and \hat{s}_2 that are paired with states s_1 and s_2), then $n_{s_2}^{\text{ind}}$ also included the number of occurrences of a transition from \hat{s}_1 to \hat{s}_2 when $u_t = \text{Ind}$. For a state in a dependent context, \hat{s}_1 , the transition count $n_{\hat{s}_2}^{\text{dep}}$ included transitions from \hat{s}_1 to \hat{s}_2 when $u_t = \text{Dep}$.

The observation parameters for a state in an independent context were sampled using the following Gibbs moves:

$$p(\mu_s | \dots) \propto N\left(\left(\sigma_s^{-2} \sum_{t \in \mathcal{I}_s} y_t + \kappa_1 \xi\right) \sigma_{\mu}^2, \sigma_{\mu}^2\right) \quad (\text{A.8})$$

$$p(\sigma_s^{-2} | \dots) \propto \Gamma\left(\alpha_1 + \frac{1}{2} |\mathcal{I}_s|, \beta_1 + \frac{1}{2} \sum_{t \in \mathcal{I}_s} (y_t - \mu_s)\right), \quad (\text{A.9})$$

where $\sigma_\mu^{-2} = \sigma_s^{-2}|\mathcal{T}_s| + \kappa_1$ and \mathcal{T}_s is the set of times for which $S_t = s$ or for which $v_t = \text{Ind}$ and $S_t = \hat{s}$. The observation parameters for a state in a dependent context were sampled using the following Gibbs moves

$$p(\mu_{\hat{s}} | \dots) \propto N\left(\left(\sigma_{\hat{s}}^{-2} \sum_{t \in \mathcal{T}_{\hat{s}}} y_t + \kappa_2(\mu_s + h)\right) \sigma_\mu^2, \sigma_\mu^2\right) \quad (\text{A.10})$$

$$p(\sigma_{\hat{s}}^{-2} | \dots) \propto \Gamma\left(\alpha_2 + \frac{1}{2}|\mathcal{T}_{\hat{s}}|, \alpha_2 \sigma_s^2 + \frac{1}{2} \sum_{t \in \mathcal{T}_{\hat{s}}} (y_t - \mu_{\hat{s}})^2\right), \quad (\text{A.11})$$

where $\sigma_\mu^{-2} = \sigma_s^{-2}|\mathcal{T}_s| + \kappa_2$ and \mathcal{T}_s is the set of times for which $v_t = \text{Dep}$ and $S_t = \hat{s}$.

Five thousand samples from the importance distribution were used to estimate equation 2.6. The availability of samples from the posterior distribution permitted evaluation of equation 2.6 with the computationally more expensive technique of bridge sampling. Unlike traditional importance sampling, an accurate estimate via bridge sampling does not require the importance distribution to be broader than the posterior distribution (Meng & Wong, 1996; Frühwirth-Schnatter, 2004). In our many test cases, we found that the differences based on which integration technique was used were negligible, suggesting that the importance distribution was accurately covering the entire mass of the posterior distribution.

Acknowledgments

This work was supported by NIH grant MH59932 and the Pennsylvania Tobacco Settlement Fund.

References

- Abraham, W. C., Logan, B., Greenwood, J. M., & Dragunow, M. (2002). Induction and experience-dependent consolidation of stable long-term potentiation lasting months in the hippocampus. *Journal of Neuroscience*, 22(21), 9626–9634.
- Agster, K. L., Fortin, N. J., & Eichenbaum, H. (2002). The hippocampus and disambiguation of overlapping sequences. *Journal of Neuroscience*, 22(13), 5760–5768.
- Ainge, J. A., van der Meer, M. A. A., & Wood, E. R. (2005). Disparity between sequence-dependent hippocampal activity and hippocampal lesion effects on a continuous t-maze task. Program no. 776.13. In *Society for Neuroscience Abstracts*, Washington, DC: Society for Neuroscience.
- Ainge, J. A., & Wood, E. (2003). Excitotoxic lesions of the hippocampus impair performance on a continuous alternation t-maze task with short delays but not with no delay. Program no. 91.1. In *Society for Neuroscience Abstracts*. Washington, DC: Society for Neuroscience.

- Barnes, C. A., Suster, M. S., Shen, J., & McNaughton, B. L. (1997). Multistability of cognitive maps in the hippocampus of old rats. *Nature*, *388*(6639), 272–275.
- Berger, T. W., & Orr, W. B. (1983). Hippocampectomy selectively disrupts discrimination reversal conditioning of the rabbit nictitating membrane response. *Behav. Brain Res.*, *8*(1), 49–68.
- Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*(2), 193–206.
- Bower, M. R., Euston, D. R., & McNaughton, B. L. (2005). Sequential-context-dependent hippocampal activity is not necessary to learn sequences with repeated elements. *J. Neurosci.*, *25*(6), 1313–1323.
- Brown, J. E., & Skaggs, W. E. (2002). Discordant coding of spatial locations in the rat hippocampus. *Journal of Neurophysiology*, *88*, 1605–1613.
- Brunswick, E. (1939) Probability as a determiner of rat behavior. *J. Exp. Psychol.*, *25*, 175–197.
- Buytendijk, F. J. J. (1930). Über das umlernen. *Arch. Néerl. Physiol.*, *15*, 283–310.
- Capaldi, E. J. (1994). The sequential view: From rapidly fading stimulus traces to the organization of memory and the abstract concept of number. *Psychonomic Bulletin and Review*, *1*, 156–181.
- Capaldi, E. J., & Neath, I. (1995). Remembering and forgetting as context discrimination. *Learning and Memory*, *2*, 107–132.
- Chib, S. (1996). Calculating posterior distributions and modal estimates in Markov mixture models. *Journal of Econometrics*, *75*, 79–97.
- Chiba, A. A., Kesner, R. P., & Reynolds, A. M. (1994). Memory for spatial location as a function of temporal lag in rats: Role of hippocampus and medial prefrontal cortex. *Behavioral and Neural Biology*, *61*, 123–131.
- Courville, A. C. (2006). *A latent cause theory of classical conditioning*. Unpublished doctoral dissertation, Carnegie Mellon University.
- Courville, A. C., Daw, N. D., Gordon, G. J., & Touretzky, D. S. (2003). Model uncertainty in classical conditioning. In S. Thrun, L. Saul, & B. Schölkopf (Eds.), *Advances in neural information processing systems*, *16* (pp. 977–984). Cambridge, MA: MIT Press.
- Courville, A. C., Daw, N. D., & Touretzky, D. S. (2004). Similarity and discrimination in classical conditioning: A latent variable account. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems*, *17* (pp. 313–320). Cambridge, MA: MIT Press.
- Daw, N. D., Courville, A. C., & Touretzky, D. S. (2006). Representation and timing in theories of the dopamine system. *Neural Comput.*, *18*(7), 1637–1677.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.*, *8*(12), 1704–1711.
- Doboli, S., Minai, A. A., & Best, P. J. (2000). Latent attractors: A model for context-dependent place representations in the hippocampus. *Neural Computation*, *12*(5), 1003–1037.
- Dufort, R. H., Guttman, N., & Kimble, G. A. (1954). One-trial discrimination reversal in the white rat. *J. Comp. Physiol. Psychol.*, *47*(3), 248–249.
- Elam, C. B., & Tyler, D. W. (1958). Reversal-learning following partial reinforcement. *Am. J. Psychol.*, *71*(3), 583–586.

- Feng, R., Rampon, C., Tang, Y., Shrom, D., Jin, J., Kiyin, M., et al. (2001). Deficient neurogenesis in forebrain-specific *presenilin-1* knockout mice is associated with reduced clearance of hippocampal memory traces. *Neuron*, *32*, 911–926.
- Ferbinteanu, J., & Shapiro, M. L. (2003). Prospective and retrospective memory coding in the hippocampus. *Neuron*, *40*(6), 1227–1239.
- Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*, *5*(5), 458–462.
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*, 680–683.
- Frank, L. M., Brown, E. N., & Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*, *27*, 169–178.
- Frühwirth-Schnatter, S. (2004). Estimating marginal likelihoods for mixture and Markov switching models using bridge sampling techniques. *Econometrics Journal*, *7*, 143–167.
- Fuhs, M. C., & Touretzky, D. S. (2000). Synaptic learning models of map separation in hippocampus. *Neurocomputing*, *32*, 379–384.
- Fuhs, M. C., & Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *J. Neurosci.*, *26*(16), 4266–4276.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M. B. (2004). Spatial representation in the entorhinal cortex. *Science*, *305*, 1258–1264.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychol. Rev.*, *107*(2), 289–344.
- Gilbert, P. E., Kesner, R. P., & Lee, I. (2001). Dissociating hippocampal subregions: A double dissociation between dentate gyrus and CA1. *Hippocampus*, *11*, 626–636.
- Gluck, M. A., & Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, *3*(4), 491–516.
- Gluck, M. A., & Myers, C. E. (1996). Integrating behavioral and physiological models of hippocampal function. *Hippocampus*, *6*(6), 643–653.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognit. Psychol.*, *51*(4), 334–384.
- Grosslight, J. H., Hall, J. F., & Scott, W. (1954). Reinforcement schedules in habit reversal: A confirmation. *J. Exp. Psychol.*, *48*(3), 173–174.
- Guzowski, J. F., Knierim, J. J., & Moser, E. I. (2004). Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron*, *44*(4), 581–584.
- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806.
- Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput.*, *14*(4), 793–817.
- Hasselmo, M. E., & Eichenbaum, H. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Netw.*, *18*(9), 1172–1190.
- Hasselmo, M. E., & Schnell, E. (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: Computational modeling and brain slice physiology. *Journal of Neuroscience*, *14*(6), 3898–3914.

- Hasselmo, M. E., Schnell, E., & Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *J. Neurosci.*, *15*(7 Pt 2), 5249–5262.
- Hasselmo, M. E., Wyble, B. P., & Wallenstein, G. V. (1996). Retrieval of episodic memories: Role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus*, *6*(6), 693–708.
- Hayman, R. M. A., Chakraborty, S., Anderson, M. I., & Jeffery, K. J. (2003). Context-specific acquisition of location discrimination by hippocampal place cells. *European Journal of Neuroscience*, *18*, 1–10.
- Hirsh, R. (1974). The hippocampus and contextual retrieval of information from memory: A theory. *Behavioral Biology*, *12*, 421–444.
- Hok, V., Save, E., Lenck-Santini, P. P., & Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proc. Natl. Acad. Sci. U.S.A.*, *102*(12), 4602–4607.
- Hölscher, C., Jacob, W., & Mallot, H. A. (2004). Learned association of allocentric and egocentric information in the hippocampus. *Exp. Brain Res.*, *158*(2), 233–240.
- Howard, M. W., Fotedar, M. S., Datey, A. V., & Hasselmo, M. E. (2005). The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychol. Rev.*, *112*(1), 75–116.
- Jarrard, L. E. (1993). On the role of the hippocampus in learning and memory in the rat. *Behavioral Neural Biology*, *60*(1), 9–26.
- Jeffery, K. J. (2000). Plasticity of the hippocampal cellular representation of space. In C. Holscher (Ed.), *Neuronal mechanisms of memory formation* (pp. 100–124). Cambridge: Cambridge University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*, 773–795.
- Kentros, C., Hargreaves, E., Hawkins, R. D., Kandel, E. R., Shapiro, M., & Muller, R. U. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science*, *280*, 2121–2126.
- Kesner, R. P., Gilbert, P. E., & Barua, L. A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behavioral Neuroscience*, *116*(2), 286–290.
- Kesner, R. P., & Novak, J. M. (1982). Serial position curve in rats: Role of the dorsal hippocampus. *Science*, *218*(4568), 173–175.
- Kimble, D. P., & Kimble, R. J. (1965). Hippocampectomy and response perseveration in the rat. *J. Comp. Physiol. Psychol.*, *60*(3), 474–476.
- Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1995). Place cells, head direction cells, and the learning of landmark stability. *Journal of Neuroscience*, *15*, 1648–1659.
- Lee, I., & Kesner, R. P. (2002). Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. *Nature Neuroscience*, *5*(2), 162–168.
- Lee, I., & Kesner, R. P. (2003). Differential roles of dorsal hippocampal subregions in spatial working memory with short versus intermediate delay. *Behavioral Neuroscience*, *117*(5), 1044–1053.

- Lee, I., Yoganarasimha, D., Rao, G., & Knierim, J. J. (2004). Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature*, *430*, 456–459.
- Lenck-Santini, P. P., Save, E., & Poucet, B. (2001). Place-cell firing does not depend on the direction of turn in a Y-maze alternation task. *Eur. J. Neurosci.*, *13*(5), 1055–1058.
- Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2006). Expansion recoding is expressed in CA3 but not in the dentate gyrus. Program no. 68.3. In *Society for Neuroscience Abstracts*. Washington, DC: Society for Neuroscience.
- Leutgeb, J. K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C. A., McNaughton, B. L., et al. (2005). Progressive transformation of hippocampal neuronal representations in “morphed” environments. *Neuron*, *48*(2), 345–358.
- Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, *309*(5734), 619–623.
- Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*, *305*(5688), 1295–1298.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., & O’Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, *416*, 90–94.
- Levy, W. B. (1989). A computational approach to hippocampal function. In R. D. Hawkins & G. H. Bower (Eds.), *Computational models of learning in simple neural systems* (pp. 243–305). San Diego, CA: Academic Press.
- Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, *6*(6), 579–591.
- Lörincz, A., & Buzsáki, G. (2000). Two-phase computational model training long-term memories in the entorhinal-hippocampal region. *Ann. N.Y. Acad. Sci.*, *911*, 83–111.
- Markus, E. J., Qin, Y., Leonard, B., Skaggs, W. E., McNaughton, B. L., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, *15*, 7079–7094.
- Marr, D. (1971). Simple memory: A theory of archicortex. *Philosophical Transactions of the Royal Society of London*, *262*(841), 23–81.
- McDonald, R. J., Ko, C. H., & Hong, N. S. (2002). Attenuation of context-specific inhibition on reversal learning of a stimulus-response task in rats with neurotoxic hippocampal damage. *Behav. Brain Res.*, *136*(1), 113–126.
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the “cognitive map.” *Nat. Rev. Neurosci.*, *7*(8), 663–678.
- McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, *10*(10), 408–415.
- Meng, X.-L. & Wong, W. A. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica*, *6*, 831–860.

- Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, *7*, 1951–1968.
- Myers, C. E., Gluck, M. A., & Granger, R. (1995). Dissociation of hippocampal and entorhinal function in associative learning: A computational approach. *Psychobiology*, *23*, 116–138.
- Nadel, L., & Willner, J. (1980). Context and conditioning: A place for space. *Physiological Psychology*, *8*(2), 218–228.
- Nakazawa, K., Quirk, M. C., Chitwood, R. A., Watanabe, M., Yeckel, M. F., Sun, L. D., et al. (2002). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science*, *297*, 211–218.
- Neave, N., Lloyd, S., Sahgal, A., & Aggleton, J. P. (1994). Lack of effect of lesions in the anterior cingulate cortex and retrosplenial cortex on certain tests of spatial memory in the rat. *Behavioural Brain Research*, *65*, 89–101.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, *34*, 171–175.
- O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, *4*(6), 661–682.
- Otmakhova, N. A., Lewey, J., Asrican, B., & Lisman, J. E. (2005). Inhibition of perforant path input to the CA1 region by serotonin and noradrenaline. *J. Neurophysiol.*, *94*(2), 1413–1422.
- Otmakhova, N. A., & Lisman, J. E. (1998). Dopamine selectively inhibits the direct cortical pathway to the CA1 hippocampal region. *J. Neurosci.*, *19*(4), 1437–1445.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev.*, *87*(6), 532–552.
- Pennes, E. S., & Ison, J. R. (1967). Effects of partial reinforcement on discrimination learning and subsequent reversal or extinction. *J. Exp. Psychol.*, *74*(2), 219–224.
- Pubols, B. H. (1962). Serial reversal learning as a function of the number of trails per reversal. *Journal of Comparative and Physiological Psychology*, *55*(1), 66–68.
- Quirk, G. J., Muller, R. U., Kubie, J. L., & Ranck Jr., J. B. (1992). The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, *12*(5), 1945–1963.
- Redish, A. D. (1999). *Beyond the cognitive map: From place cells to episodic memory*. Cambridge, MA: MIT Press.
- Redish, A. D. (2001). The hippocampal debate: Are we asking the right questions? *Behavioral Brain Research*, *127*, 81–98.
- Redish, A. D., & Touretzky, D. S. (1998). The role of the hippocampus in solving the Morris water maze. *Neural Computation*, *10*(1), 73–111.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokesy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Robitsek, R. J., Fortin, N., & Eichenbaum, H. (2005). Hippocampal unit activity during continuous and delayed t-maze spatial alternation. Program no. 776.7. In *Society for Neuroscience Abstracts*, Washington, DC: Society for Neuroscience.

- Rolls, E. T. (1996). A theory of hippocampal function in memory. *Hippocampus*, 6, 601–620.
- Rudy, J. W., & Sutherland, R. J. (1995). Configural association theory and the hippocampal formation: An appraisal and reconfiguration. *Hippocampus*, 5, 375–389.
- Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17(15), 5900–5920.
- Schmajuk, N. A., & DiCarlo, J. J. (1991). A neural network approach to hippocampal function in classical conditioning. *Behav. Neurosci.*, 105(1), 82–110.
- Schmajuk, N. A., & DiCarlo, J. J. (1992). Stimulus configuration, classical conditioning, and hippocampal function. *Psychol. Rev.*, 99(2), 268–305.
- Schmitz, D., Mellor, J., Breustedt, J., & Nicoll, R. A. (2003). Presynaptic kainate receptors impart an associative property to hippocampal mossy fiber long-term potentiation. *Nature Neuroscience*, 6(10), 1058–1063.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6, 461–464.
- Shapiro, M. L., Tanila, H., & Eichenbaum, H. (1997). Cues the hippocampal place cells encode: Dynamic and hierarchical representation of local and distal stimuli. *Hippocampus*, 7(6), 624–642.
- Silveira, J. M., & Kimble, D. P. (1968). Brightness discrimination and reversal in hippocampally-lesioned rats. *Physiology and Behavior*, 3, 625–630.
- Skaggs, W. E., & McNaughton, B. L. (1998). Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *Journal of Neuroscience*, 18(20), 8455–8466.
- Smith, D. M., & Mizumori, S. J. Y. (2006). Learning-related development of context-specific neuronal responses to places and events: The hippocampal role in context processing. *J Neurosci*, 26(12), 3154–3163.
- Tabuchi, E., Mulder, A. B., & Wiener, S. I. (2003). Reward value invariant place responses and reward site associated activity in hippocampal neurons of behaving rats. *Hippocampus*, 13(1), 117–132.
- Tanila, H., Shapiro, M. L., & Eichenbaum, H. (1997). Discordance of spatial representation in ensembles of hippocampal place cells. *Hippocampus*, 7(6), 613–623.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.*, 24(4), 629–640, 652–791.
- Touretzky, D. S., & Redish, A. D. (1996). A theory of rodent navigation based on interacting representations of space. *Hippocampus*, 6(3), 247–270.
- Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, 2, 189–199.
- Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of hippocampus in memory. *Hippocampus*, 4(3), 374–391.
- Tsodyks, M. (1999). Attractor neural network models of spatial maps in hippocampus. *Hippocampus*, 9, 481–489.
- Vazdarjanova, A., & Guzowski, J. F. (2004). Differences in hippocampal neuronal population responses to modifications of an environmental context: Evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. *J. Neurosci.*, 24(29), 6489–6496.

- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, *13*, 260–269.
- Wallenstein, G. V., Eichenbaum, H., & Hasselmo, M. E. (1998). The hippocampus as an associator of discontinuous events. *Trends in Neuroscience*, *21*(8), 317–323.
- Wallenstein, G. V., & Hasselmo, M. E. (1998). GABAergic modulation of hippocampal population activity: Sequence learning, place field development, and the phase precession effect. *Journal of Neurophysiology*, *78*, 393–408.
- Wike, E. L. (1953). Extinction of a partially and continuously reinforced response with and without a rewarded alternative. *J. Exp. Psychol.*, *46*(4), 255–260.
- Winocur, G., & Olds, J. (1978). Effects of context manipulation on memory and reversal learning in rats with hippocampal lesions. *J. Comp. Physiol. Psychol.*, *92*(2), 312–321.
- Wise, L. M. (1962). Supplementary report: The Weinstock partial reinforcement effect and habit reversal. *J. Exp. Psychol.*, *64*, 647–648.
- Wood, E. R., Dudchenko, P. A., & Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, *397*, 613–616.
- Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000) Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, *27*, 623–633.
- Xavier, G. F., Oliveira-Filho, F. J. B., & Santos, A. M. G. (1999). Dentate gyrus-selective colchicine lesion and disruption of performance in spatial tasks: Difficulties in “place strategy” because of a lack of flexibility in the use of environmental cues. *Hippocampus*, *9*, 668–681.
- Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, *46*(4), 681–692.