nature neuroscience



# Matching storage and recall: hippocampal spike timing–dependent plasticity and phase response curves

Máté Lengyel[1], Jeehyun Kwag[2], Ole Paulsen[2] & Peter Dayan[1]

**Hippocampal area CA3 is widely considered to function as an autoassociative memory. However, it is insufficiently understood how it does so. In particular, the extensive experimental evidence for the importance of carefully regulated spiking times poses the question as to how spike timing–based dynamics may support memory functions. Here, we develop a normative theory of autoassociative memory encompassing such network dynamics. Our theory specifies the way that the synaptic plasticity rule of a memory constrains the form of neuronal interactions that will retrieve memories optimally. If memories are stored by spike timing–dependent plasticity, neuronal interactions should be formalized in terms of a phase response curve, indicating the effect of presynaptic spikes on the timing of postsynaptic spikes. We show through simulation that such memories are competent analog autoassociators and demonstrate directly that the attributes of phase response curves of CA3 pyramidal cells recorded *in vitro* qualitatively conform with the theory.**

The task of storing memories and recalling them from partial or noisy cues is fundamental for the brain and has been a particular focus for empirical[1,2] and theoretical work[3] on the hippocampus. This naturally raises the key question as to how the properties of single cells and the overall hippocampal network support its proposed function. The CA3 region of the hippocampus has the densest recurrent collateral system in the mammalian cortex[4], which is consistent with the central role accorded to recurrent connections in standard models of autoassociative memories[5]. However, with relatively few exceptions (in the hippocampus and elsewhere[6–13]), such models typically use a highly simplified treatment of the resulting collective dynamics of their model neurons[14,15]. They thereby fail to capture a salient characteristic of the activity of hippocampal neurons during memory states: namely, the central role played by spike timing.

Evidence for the importance of timing in the hippocampus is widespread. Behaviorally relevant neural oscillations at different frequencies pace the activity of all hippocampal cell types. In addition, the timing of individual action potentials of principal cells is tightly regulated[16,17], and temporal sequences of the ensemble firing pattern consistently reappear during both awake behavior[18,19] and sleep[20–22]. Furthermore, synaptic plasticity is also critically sensitive to the precise timing of pre- and postsynaptic spikes[23,24].

A key idea for understanding networks such as CA3 in which information may be coded by spike times is the phase response curve (PRC). PRCs offer a precise characterization of the effect a presynaptic spike has on the timing of the succeeding postsynaptic spike depending on its time of arrival[25–27]. PRCs and related phase reduction and analysis methods have wide application in everything from the analysis of cardiac rhythms[28] and patterns of oscillatory coordination for motor pattern generation[29] to the relationship between nonlinear and subthreshold intrinsic mechanisms within cells and various forms of synchrony[30,31]. However, to our knowledge, hitherto they have not been used to characterize oscillatory autoassociative memories.

Here, we present a theory treating autoassociative recall as optimal probabilistic inference[32,33], inferring the recurrent dynamics within a memory that are normatively matched to the form of the synaptic plasticity rule used to store traces. In the case of CA3, this encompasses memories encoded in spike timings relative to underlying neural oscillations, and it thus involves inferring the optimal PRC of neurons from their spike timing–dependent synaptic plasticity (STDP) rule. We thus make specific predictions about the shape of the PRC of CA3 pyramidal neurons based on the STDP reported in cultured hippocampal neurons[34]. We show that the theoretically derived PRC provides a good qualitative match to those recorded in hippocampal CA3 pyramidal cells *in vitro*.

## RESULTS

### Autoassociative recall as probabilistic inference

The fundamental requirement for an autoassociative memory is to recall a previously encoded memory trace when cued with a noisy or partial cue. As synaptic plasticity, the key mechanism for long term storage in the brain, loses information about the traces, recall poses a complex problem of probabilistic inference. The optimal solution to this problem, which amounts to a normative theory of recall, depends

[1]Gatsby Computational Neuroscience Unit, University College London, 17 Queen Square, London WC1N 3AR, UK. [2]University Laboratory of Physiology, Oxford University, Parks Road, Oxford OX1 3PT, UK. Correspondence should be addressed to M.L. (lmate@gatsby.ucl.ac.uk).

Received 19 April; accepted 12 September; published online 30 October 2005; doi:10.1038/nn1561

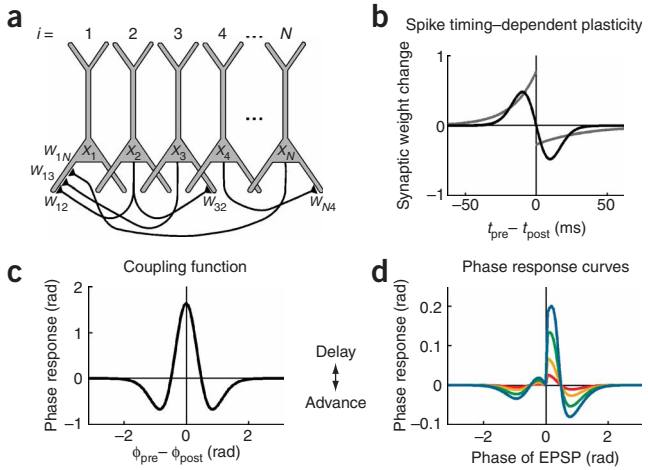


**Figure 1** Normative theory of spike timing–based autoassociative memory. (**a**) Schematic diagram of a recurrent network of neurons. Neurons are numbered $i = 1 \ldots N$ and are characterized by their respective activities, $x_1 \ldots x_N$. Presynaptic neuron $j$ is connected to postsynaptic neuron $i$ through a recurrent synapse with efficacy (weight) $w_{ij}$. Although all-to-all connectivity (excluding autapses) was assumed for the theoretical derivations, here only a few synapses are shown for clarity. (**b**) Memories are stored by a spike timing-dependent plasticity (STDP) rule derived from experiments on cultured hippocampal neurons[34]. $t_{pre}$ and $t_{post}$ represent times of pre- and postsynaptic firing. Gray lines are exponential fits[24] to data from ref. 34. Black line is a continuous fit taken to be the synaptic learning rule ($\Omega$) in equation (1). (**c**) Optimal coupling function (H) for retrieving memories stored by STDP (black line in **b**), as derived in equation (3). This shows how the firing phase of the postsynaptic neuron should change as a function of the phase difference between pre- and postsynaptic firing, if neurons were to interact continuously. $\phi_{pre}$ and $\phi_{post}$ represent firing phases of pre- and postsynaptic cells relative to a local field potential oscillation. (**d**) Optimal phase response curves (PRCs) derived from the optimal coupling function (shown in **c**), showing how neurons should interact through spikes. Different curves correspond to linearly increasing synaptic weights (in increasing order: red, yellow, green, blue). 'Zero' phase is the phase of the postsynaptic spike.

on the statistical characteristics of the traces and cues and the nature of the lossiness of storage.

More concretely, consider a network of $N$ neurons (**Fig. 1a**), fully connected by $N \times (N - 1)$ synaptic connections, and storing $M$ memory traces. Memories are represented by distributed patterns of neural activity, with scalar variable $x_i^m$ characterizing the $i$th neuron in the $m$th memory trace. Here, we treat $x_i^m$ as neuron $i$'s spiking time relative to a reference point of an ongoing field potential oscillation, such as the peak of theta oscillation[16,35]. Storage amounts to changing the synaptic weights between neurons according to their activities in a memory trace using a synaptic plasticity rule:

$$\Delta w_{ij} = \Omega(x_i^m, x_j^m) \qquad (1)$$

This rule is local in that the change to the synaptic weight $w_{ij}$ between presynaptic neuron $j$ and postsynaptic neuron $i$ depends only on the activities of these two neurons and not those of other neurons. Except for this constraint, we allow $\Omega$ to be an arbitrary function. We also make the simplifying assumption that synaptic plasticity is additive across the memories:

$$w_{ij} = \sum_m \Omega(x_i^m, x_j^m) \qquad (2)$$

Local and additive plasticity loses information in storing memories, because $\Omega$ transforms pre- and postsynaptic activity into a single scalar contribution. Indeed, equation (2) is noninvertible, with many different combinations of synaptic weight changes potentially leading to the same total synaptic weight. This lossiness implies that recall from a noisy or partial cue involves a process of combining probabilistic information from (i) the general statistical properties of the traces (that is, the prior distribution), (ii) the cue itself and (iii) the synaptic weights. Finding the statistically most likely memory trace involves complex, nonlocal operations. However, it can be well approximated (**Supplementary Note** online) by a form of neural dynamics among the recurrently coupled neurons in which there are explicit terms corresponding exactly to each of these three sources of information.

Under this account, synaptic interactions between neurons of the network implement a 'search' for the activity pattern associated with the original memory trace that was most likely to have led to the cue. Over the course of search, each neuron gradually changes its spike timing such that the activities of it and its synaptic partners are increasingly likely to reflect a pattern consistent with the corresponding synaptic weights. As a result, the contribution to the dynamics of neuron $i$ associated with the synaptic weight term involves a linear sum of influences from its presynaptic afferents $j$, with each element in the sum taking the form

$$\mathrm{H}(x_i, x_j) = w_{ij} \frac{\partial}{\partial x_i} \Omega(x_i, x_j) \qquad (3)$$

This has the obvious appealing characteristic that the strength of the interaction between two neurons is scaled by the synaptic weight. Less obvious is our key suggestion that the interaction should be determined by the derivative of the synaptic plasticity rule that was used to store memories in the network. We can derive an intuition about this rule by considering what happens if the synaptic weight is positive; that is, above an average value. This large synaptic weight arises from the weight changes associated with the memory traces. Therefore, the neuron should change its activity to increase the contribution that it and its synaptic partner would have made to the synaptic weight had their present activity indeed been associated with one of the memories that caused this excess synaptic weight.

Here, we study the case of area CA3 in the hippocampus, in which the synaptic plasticity rule involves STDP. We show that the optimal dynamics for neurons representing memory traces in terms of the phase of firing relative to an underlying oscillation is determined by a particular shape of PRC that we experimentally validate.

## Spike timing-based memories

The hippocampus, as well as other areas involved in memory processing, demonstrates prominent local field potential oscillations (LFPOs) under a variety of conditions, including both awake and sleep states[36]. In such cases, the phases of the spikes of a neuron relative to the LFPO have been shown to be carefully controlled[19] and even to convey meaningful stimulus information, such as information about the position of the animal in its environment[16]. The discovery of STDP, for which the relative timing of pre- and postsynaptic firings determines the sign and extent of synaptic weight change, offered new insights into how the information represented by spike times might be stored[37]. However, except for some interesting suggestions about neuronal resonance[12], it is less clear how one might correctly recall this information.

Our theory allows a systematic treatment of this case, if we interpret neuronal activities as the phases of firing relative to an ongoing LFPO. This description is valid in the limit that neurons are driven to spike by a perithreshold oscillation approximately once in each cycle of the LFPO. First, we interpret storage ($\Omega$ in equation (1)) in terms of an STDP rule (**Fig. 1b**; recorded in cultured hippocampal neurons[34]), with

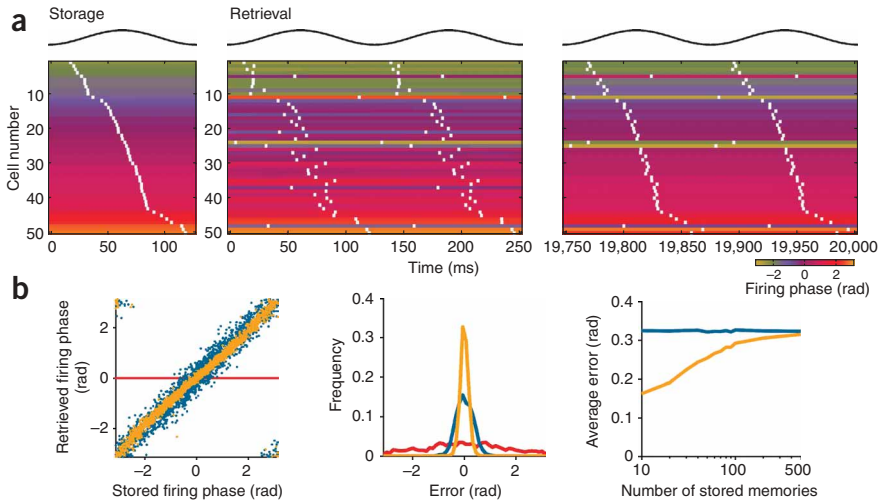


**Figure 2** Quality of memory retrieval in an optimally constructed spike timing–based autoassociative memory model: numerical simulations. (**a**) Retrieval of a memory pattern defined by firing phases. Left: noise-free form of the cued pattern as originally stored. Each row shows the firing phase of a cell (color code) and its notional spike time (white squares) depending on when the phase of the underlying theta oscillation (top trace) coincided with the firing phase of the cell. Cells in all three panels were ordered according to their firing phases in the noise-free memory trace. One complete cycle of theta oscillation (125 ms) is shown. Center: first two theta cycles during retrieval. Right: two theta cycles from the end of the retrieval process. (**b**) Retrieval performance of the network. Three networks were compared: a 'prior only' network (red), an 'input only' network (blue) and the complete optimal network (yellow; see text for further details). Left: retrieved firing phases (*y*-axis) as a function of noise-free firing phases of the associated cued traces (*x*-axis). Points near the diagonal indicate good retrieval. Center: histogram of errors (circular difference of retrieved and stored firing phases) for the three networks. Right: root mean squared error over neurons for the 'input only' and the complete network as a function of the number of memories stored. The 'prior only' network is omitted from this plot for clarity, because its average error was close to the maximally possible $\pi/2$ value.



the relative phases translated directly into relative spiking times assuming a theta-frequency oscillation. This rule prescribes that a synapse is strengthened if the presynaptic neuron spikes before the postsynaptic neuron and is weakened if the order of spiking is reversed. The data around zero time difference between pre- and postsynaptic firing are variable, so we consider a smooth, differentiable fit (**Fig. 1b**, black line) that captures the salient characteristics of hippocampal STDP. Second, given this firing phase interpretation of storage, the optimal recall dynamics of equation (3) also acts on firing times, requiring the phases of the spikes of a postsynaptic cell to be advanced or delayed relative to the LFPO on the basis of the phases at which the neuron itself and its presynaptic partners spike.

The STDP rule specifies weight changes based on the difference between pre- and postsynaptic spike times. Thus, in recall, the interaction between neurons should also be a function of this difference, sometimes called a phase coupling function[38]. Specifically, equation (3) tells us that the STDP rule shown in **Figure 1b** requires the phase coupling function shown in **Figure 1c**. However, since influences must be based on discrete spikes, a correction is necessary (**Supplementary Note**). Interactions are then described in terms of a PRC[25] (**Fig. 1d**), which indicates the extent to which the timing of the next spike of a cell is advanced or delayed as a function of the timing relative to its most recent spike of a small perturbation caused, for instance, by an excitatory postsynaptic potential (EPSP). This relationship between the PRC of a neuron and the STDP rule employed by its synapses provides a way of testing our theory.

Since our derivations embody assumptions and approximations, we tested the recall performance of the network by numerical simulations (**Fig. 2**). Memories were stored using an STDP rule derived from the hippocampus (**Fig. 1b**, black line), and retrieved by a network using the optimally matching coupling function (**Fig. 1c**). During the course of retrieval of a single memory trace (shown with one whole theta cycle in **Fig. 2a**, left), firing phases smoothly changed from an initially noisy value reflecting the input to the network (**Fig. 2a**, center, first theta cycle) to the final, converged form (**Fig. 2a**, right). The noise-free version of the cue (**Fig. 2a**, left) was recovered with high fidelity. There were no discernible changes in spike times in the last two cycles, demonstrating that the network had reached an attractor state. Note that some of the noise was already removed by the second theta cycle.

The network incorporated information from the three sources of evidence that pertain to recall: the prior distribution, the input and the synaptic weights (**Fig. 2b**). We used two purely feed-forward networks as benchmarks of the optimal recurrent network to dissect the contribution of interactions through the recurrent synapses. One network used only information in the prior distribution and thus always retrieved only 0, the mean of the prior distribution, whereas the other used only information in the input and thus simply transmitted its input to its output. Compared with these networks, individual values recalled by the full network were closer to their ideal values (**Fig. 2b**, left). The distributions of errors were symmetrical around zero (correct retrieval) for all three networks, indicating that all of them were unbiased (**Fig. 2b**, middle). However, the complete network was markedly superior over the other two in terms of variance of the retrieved firing phases around the correct values; that is, it made smaller errors on average (**Fig. 2b**, right). The comparison with using only the input is especially important, as persistent input is known to improve recall performance by itself[39]. Of course, performance did ultimately deteriorate as an increasing number of patterns was stored, and information in synaptic weights became negligible relative to information in the noisy input (**Fig. 2b**, right). This graceful degradation arose from optimal integration of available information sources (synaptic weights and input, in this case).

We also tested the robustness of our results in a number of adversarial settings (**Supplementary Fig. 1** online), including storage noise and limited connectivity. Recall performance was proportional to the degree of connectivity and inversely proportional to storage noise (**Supplementary Fig. 1**). Performance was also proportional to the size of the network (**Supplementary Fig. 1**), as is common for such memories[3]. Although both the network sizes (50–250 neurons) and connectivity ratios (100–20%) used in our simulations were outside the range of realistic values for rat CA3 (300,000 neurons with 5% connectivity[4]), these results together imply that the main determinant of recall performance is the number of synapses per neuron, and thus we predict that a network with realistic anatomy

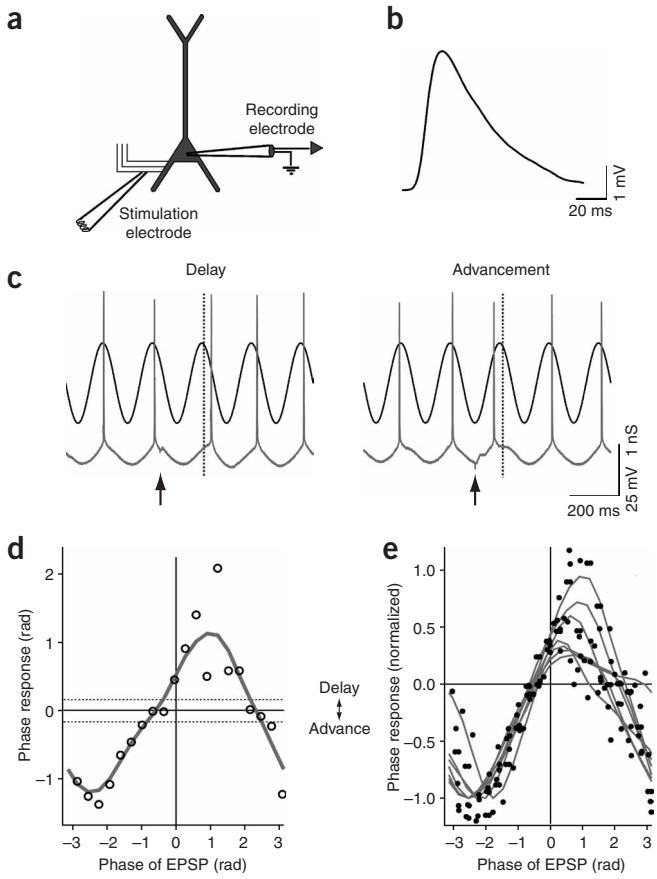


**Figure 3** Experimental measurement of the PRC in CA3 hippocampal neurons. (**a**) Diagram of a CA3 hippocampal neuron with patch recording electrode at the soma and extracellular stimulation electrode among recurrent fibers in the stratum oriens. Sinusoidal inhibitory conductance mimicking hippocampal theta oscillation (5 Hz) was injected through the patch pipette using dynamic clamp. An EPSP was evoked using extracellular stimulation. (**b**) Average somatic EPSP recorded in response to extracellular stimulation without oscillation ($n = 5$). (**c**) Sample of current-clamp recordings showing the phase response of a CA3 neuron (gray trace) to the stimulated EPSP (arrows; times of stimulation) during 5 Hz oscillation (black trace). (**d**) Plot of phase delay and advancement of the spike as a function of the phase of the EPSP. 'Zero' phase was defined as the average phase at which spikes occurred during 5 Hz oscillation without EPSP (vertical dotted lines in **c**). The PRC (open circles) was subject to Gaussian smoothing (gray line). Horizontal dotted lines show ± 2 s.d. of the average spike phase without EPSP. (**e**) Smoothed PRCs (gray lines) and raw data points (filled black circles) normalized for $n = 7$ cells. Note that there are virtually no data points in the second quadrant.

(12,000 recurrent synapses per CA3 pyramidal cell[4], as opposed to the maximal 200 synapses used in our simulations) will be a highly competent spike timing–based autoassociator. Further, performance degraded only weakly even if the STDP was asymmetric: for instance, with the potentiation having larger maximal amplitude and tighter time frame than depression[34] (**Supplementary Fig. 1**). Performance was more sensitive to a mismatch between storage rule and recall dynamics.

### Characteristics of the optimal phase response curve

The theoretically optimal PRC (**Fig. 1d**) for autoassociative memory recall has five salient characteristics. First, excitatory currents can cause both delay (positive parts) and advancement (negative parts) of the next spike. Second, spike delay is predicted for EPSPs that follow postsynaptic spiking. Third, EPSPs immediately preceding postsynaptic spikes should have negligible effect on postsynaptic spikes. Fourth, EPSPs before this insensitive period or after the interval where delay is predicted should result in advancement. Fifth, based on equation (3) and shown as different colored lines in **Figure 1d**, the effect of presynaptic spiking on the phase response should scale with the synaptic weight between the two cells. The optimal scaling of the PRC is not exactly linear, but its zero crossings (relative spike times for which no phase shift is predicted) should be unaffected by changing the synaptic weight. Type II oscillators, such as the Hodgkin-Huxley model, show spiking behavior that broadly complies with these criteria[38], thus suggesting that real neurons may implement similar PRCs.

These features are preserved (**Supplementary Fig. 2** online) for a range of STDP curves that satisfy a few qualitative properties: potentiation for pre- before postsynaptic firings, depotentiation for post- before presynaptic firings, pre- and postsynaptic spikes required to appear within a limited time window for both, and a transitionary regime between the potentiation and depotentiation at around zero time difference.

The optimal PRC also seems to be insensitive to inputs arriving in the middle of the spiking cycle (shown as the two flat flanks of the PRC in **Fig. 1d**), unlike most biophysically plausible PRCs[26]. This insensitivity is predicted because the function used to fit experimental STDP curves converges to zero and is therefore already flat at 25 ms (**Fig. 1b**, black line), whereas for the conversion from spike time to phase, the length of the theta cycle was assumed to be 125 ms (corresponding to the widely reported 8 Hz theta frequency[35]). A shallower fall-off of the STDP curve (as shown by the original exponential fit of the data; **Fig. 1b**, gray line) or a higher theta frequency would diminish this region, leading to the fusion of the two intervals where advancement is predicted (**Supplementary Fig. 2**).

### Phase response curves of hippocampal CA3 pyramidal cells

We used somatic whole-cell patch-clamp recordings from CA3 pyramidal neurons in acute hippocampal slices to measure the PRC for comparison with the theory. Theta oscillation was simulated by a somatic oscillatory inhibitory conductance, as is also observed *in vivo*[35], and excitatory synaptic input was delivered by extracellular stimulation (**Fig. 3a,b**). We confirmed experimentally that excitatory input could both delay and advance spikes (**Fig. 3c**). As predicted by theory, delay was observed in the next cycle when EPSP followed immediately after a spike, and advancement was observed when EPSP occurred before the expected spike or well after the previous spike (**Fig. 3d,e**; $n = 7$ cells).

To confirm that both phase advancement and phase delay are due to the EPSP itself, and not to some other extracellular stimulation-evoked modulatory or network event, we repeated the experiment using dynamic clamp to simulate an excitatory input conductance (**Fig. 4**). Indeed, the same effects of phase delay and phase advancement were observed using artificial excitatory postsynaptic conductances (EPSGs, **Fig. 4a,b**). Moreover, we also confirmed that the effect on phase advancement and phase delay increased with synaptic conductance, with the zero crossings of the PRC remaining relatively unaffected (**Fig. 4c**), as predicted by theory (**Fig. 1d**). Similar results were obtained in seven other cells. Finally, in order to test the generalizability of our findings, we recorded PRCs at a higher but still within-theta band frequency (**Supplementary Fig. 3** online) as well as in response to bursts of EPSGs (**Supplementary Fig. 4** online) and found that PRCs were preserved under these conditions. In sum, individual

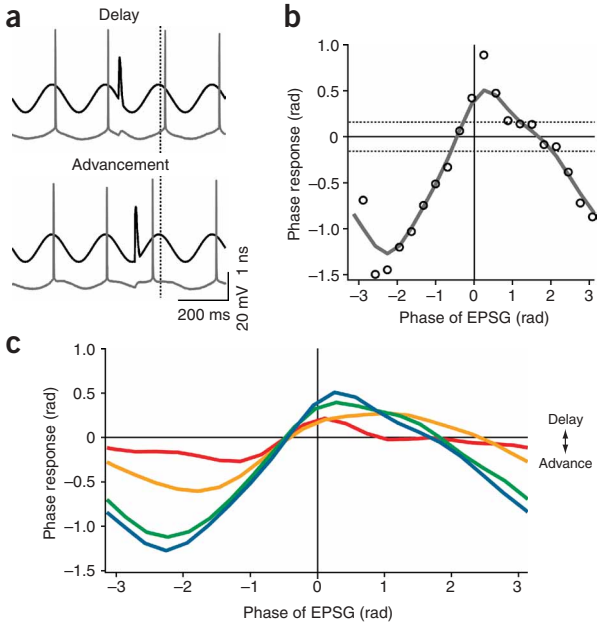


**Figure 4** Effect of EPSG amplitude on the PRC. (**a**) Sample of current-clamp recordings showing the response of a CA3 neuron (gray trace) to the EPSG during 5 Hz oscillation (black trace). The vertical dashed line represents the average phase at which spikes occurred during 5 Hz oscillation without EPSG. (**b**) Plot of phase delay and advancement of spike as a function of the phase of the EPSG (2 nS). The PRC (open circles) was subject to Gaussian smoothing (gray line). Horizontal dotted lines show ± 2 s.d. of the average spike phase without EPSG. (**c**) Smoothed PRCs produced by EPSGs of different amplitude in the same cell; 0.5 nS (red), 1 nS (yellow), 1.5 nS (green) and 2 nS (blue).

CA3 pyramidal neurons demonstrate intrinsic dynamics that support optimal retrieval of information by phase coding according to our theory.

## DISCUSSION

We report a normative theory of statistically sound recall in analog associative memory networks. We have shown that the theory makes a direct link between the rule governing spike timing–dependent synaptic plasticity and the neurons' PRCs, and we have qualitatively confirmed the precepts of the theory by recording conformant PRCs from hippocampal CA3 neurons.

Our technique treats analog autoassociative memory from a probabilistic viewpoint[32,33], deriving a general relationship between (i) the nature and representational substrate of the memory traces and the rules governing neural plasticity, and (ii) the dynamical behavior during recall that would approximately solve a formally presented task such as pattern completion or noise removal. Applied to the case of memory traces represented as phases, and stored by an STDP rule (derived from data from cultured hippocampal neurons[34]), the resulting dynamics specified a form of PRC. Not only were the general characteristics of this PRC consistent with those in the CA3 data (for instance, the existence of delays and advances), but also the more detailed predictions were matched, such as the scaling of the PRC with the input (in the dynamic clamp experiments) and even the form of the delays and advances relative to the standard firing phase. It is not at all trivial that the resulting PRC that was expected had a biophysically reasonable form, let alone that it matched actual PRCs in CA3. Indeed, in contrast to the type II–like PRCs we recorded here in hippocampal CA3 pyramidal cells, classical integrate-and-fire dynamics produce only phase advancement in response to excitatory inputs[26], and even neocortical pyramidal cells show phase response characteristics of Type I membranes and thus lack a delay component in their PRCs[40]. It is also not trivial that the network performed recall competently, as analog autoassociative memory is hard[15].

Our theoretical framework embodied a number of simplifying assumptions that allowed for an analytical derivation of the optimal recall dynamics but whose biological plausibility may seem to be unclear. We explicitly tested the incorporation of storage noise, limited connectivity and asymmetry in the STDP rule (**Supplementary Fig. 1**), showing that none of these had an importantly deleterious effect on performance. As one might expect from the framework, the most marked sensitivity is to mismatch between storage and recall (**Supplementary Fig. 1**).

One more holistic assumption was that, in line with the traditional theory of autoassociative memories[3,5], we treated memory encoding and retrieval separately, as distinct modes of operation. Specifically, neural activities during encoding were clamped to the memory patterns being stored so that the intrinsic dynamics of the network did not contribute to this process. Although, in its extreme form, this assumption is certainly unrealistic, there is suggestive data that changing levels of acetylcholine neuromodulation may result in the separation of memory encoding and retrieval in the hippocampus and related structures by selectively suppressing transmission and plasticity in afferent or internal synaptic pathways during these two operational modes[41].

Another assumption was to have addressed only the simplest form of oscillatory memory in which all neurons fired once per cycle. This was a marked abstraction of the hippocampus, whose pyramidal cells often fire bursts of action potentials *in vivo*[16,18]. The induction of synaptic plasticity is also most effective when bursts rather than single spikes are used in the stimulation protocol[23], and spike timing–dependent plasticity has been shown to encompass multi-spike interactions[42] and to be sensitive to the firing rate of pre- and postsynaptic cells[43]. Thus, an extension to a joint rate- and phase-based code for information is pressing[44–46]. We suggest that the choice of the number of spikes fired in a cycle (including no spikes) could convey orthogonal information, characterizing the certainty a neuron has about its phase or, indeed, its relevance for the given pattern (M.L. and P.D., unpublished data). Under this account, the consequences of firing potentially multiple spikes per theta cycle for memory encoding are straightforward: some memories will be stored and therefore retrieved with greater efficiency. Retrieval dynamics would also have to take into account the extra information conveyed by instantaneous firing rates. Preliminary experimental results (**Supplementary Fig. 4**) are compatible with the conclusion from the extension of our theory that interactions between bursting cells should be scaled versions of single spike-based PRCs.

An intriguing suggestion evident in the single-case figures (**Figs. 3c** and **4a**) is that, after a stimulation, not only is the phase of the very next spike altered but also the phases of a few successive spikes change[27]. Depending on the assumptions postsynaptic neuronal mechanisms might embody about a neuron's presynaptic cousins, the theory can predict various forms for these multistep PRCs; it would therefore be interesting to characterize these more fully.

Finally, oscillations in one structure are only a small part of the overall puzzle of memory. There is increasing evidence for the involvement of multiple structures that undergo oscillations of potentially different frequencies and intermittencies[36] but are nevertheless tightly and jointly regulated[10,47,48]. Perhaps a first step will be to generalize and abstract away from single-neuron PRCs to a form of population PRC,

indicating the overall effect of one oscillatory structure on another to which it has moderately dense connections and with which it is coordinated.

## METHODS

**Simulations.** Networks of $N = 200$ neurons were simulated (50 randomly chosen cells are shown in **Fig. 2a**). Memories were stored by an additive learning rule (equations (1) and (2)) that was a circular fit to experimental STDP curves[24], using a Gabor function: $\Omega(x_i,x_j) = A \exp(s \cos\Delta\phi_{ij})$, with $\Delta\phi_{ij} = (2\pi / T_\theta (x_i - x_j))$ and $T_\theta = 125$ ms. Best-fitting parameters determined by minimizing the squared error between –62.5 and 62.5 ms were $A = 0.03$ and $s = 4$ (**Fig. 1b**). The number of stored memories was $M = 10$ (**Fig. 2a**,**b**, left and center) or was varied between $M = 10$ and $M = 500$ (**Fig. 2b**, right). Firing phases in memory patterns were drawn from a von Mises (circular Gaussian) distribution with $\mu_x = 0$ mean and $k_x = 0.5$ concentration (the prior distribution), resulting in a distribution of firing phases that matched those recorded *in vivo* for hippocampal pyramidal cell populations[49,50].

At retrieval, a randomly selected pattern from the list of stored patterns was used as the recall cue corrupted with unbiased additive circular Gaussian noise of $k = 10$ concentration. Retrieval dynamics of the network was parameterized accordingly to optimally match the form and parameters of the prior and the noise distribution, as well as of the synaptic plasticity rule (**Supplementary Note**), and involved a phase coupling function $H(x_i,x_j) = 2\pi w_{ij} A / T_\theta \exp(s \cos\Delta\phi_{ij}) \cdot (\cos\Delta\phi_{ij} - s \sin^2\Delta\phi_{ij})$ (shown in **Fig. 1c**). Differential equations were solved numerically by using an adaptable step-size method, and the states of neurons were recorded every 5 ms (simulated time; **Fig. 2a**) or at the end of the simulation after 20000 ms simulated time (**Fig. 2b**). For **Figure 2b** 10 networks with different lists of stored patterns were simulated, 10 retrieval attempts were made in each network. Data points show results pooled over networks and retrieval attempts.

**Experiments.** Somatic whole-cell patch-clamp recordings were made from CA3 pyramidal cells in hippocampal slices prepared from postnatal day 13–19 Wistar rats and maintained at 29–31 °C (for details of extracellular and intracellular solutions, see **Supplementary Methods** online). Theta oscillation was simulated by 5 Hz or 8 Hz oscillatory inhibitory conductance of 1–2 nS peak amplitude using dynamic clamp. A positive tonic current was superimposed on the oscillatory input so that the membrane potential was depolarized just enough to evoke a single action potential near the peak of every cycle of the oscillation. EPSP was evoked via extracellular stimulation in the presence of 1 μM gabazine (SR95531). EPSG of peak amplitude 0.5–4.5 nS was injected using dynamic clamp at 20 different phases of the inhibitory oscillation, starting at 'zero' phase. Each PRC data point is an average of ten repetitions with the same stimulation phase. Data were acquired online and analyzed with custom-made procedures in Igor Pro. In normalizing (**Fig. 3e** and **Supplementary Figs. 3** and **4**), both smoothed PRC and raw data points were divided by the peak advancement value of the smoothed PRC for each cell.

*Note: Supplementary information is available on the Nature Neuroscience website.*


### ACKNOWLEDGMENTS

We thank B. Gutkin, D. MacKay and E. Shea-Brown for valuable discussions and E.O. Mann and D. McLelland for their help with Igor procedures. This work was supported by the Gatsby Charitable Foundation (M.L., P.D.), the European Bayesian-Inspired Brain and Artefacts project (M.L., P.D.), the Biotechnology and Biological Sciences Research Council (J.K., O.P.), the Kwanjung Educational Foundation, Korea (J.K.) and the Oxford University Clarendon Fund (J.K.).


### COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.


Published online at http://www.nature.com/natureneuroscience/
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions/



1. Squire, L.R. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* **99**, 195–231 (1992).
2. Cohen, N.J. & Eichenbaum, H. *Memory, Amnesia, and the Hippocampal System* (MIT Press, Cambridge, Massachusetts, 1993).
3. Rolls, E.T. & Treves, A. *Neural Networks and Brain Function* (Oxford Univ. Press, Oxford, 1998).
4. Amaral, D.G., Ishizuka, N. & Claiborne, B. Neurons, numbers and the hippocampal network. *Prog. Brain Res.* **83**, 1–11 (1990).
5. Hopfield, J.J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* **79**, 2554–2558 (1982).
6. Baird, B. Bifurcation and learning in network models of oscillating cortex. *Physica D.* **42**, 365–384 (1990).
7. Wang, D.L., Buhmann, J. & von der Malsburg, C. Pattern segmentation in associative memory. *Neural Comput.* **2**, 94–106 (1990).
8. Li, Z. Modeling the sensory computations of the olfactory bulb. in *Models of Neural Networks* Vol. 2 (eds. Domany, E., van Hemmen, J.L. & Schulten, K.) 221–251 (Springer-Verlag, New York, 1995).
9. Hendin, O., Horn, D. & Tsodyks, M.V. Associative memory and segmentation in an oscillatory neural model of the olfactory bulb. *J. Comput. Neurosci.* **5**, 157–169 (1998).
10. Li, Z. & Hertz, J. Odour recognition and segmentation by a model olfactory bulb and cortex. *Network* **11**, 83–102 (2000).
11. Hasselmo, M.E., Bodelon, C. & Wyble, B.P. A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput.* **14**, 793–817 (2002).
12. Scarpetta, S., Li, Z. & Hertz, J. Hebbian imprinting and retrieval in oscillatory neural networks. *Neural Comput.* **14**, 2371–2396 (2002).
13. Jensen, O. & Lisman, J. Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends Neurosci.* **28**, 67–72 (2005).
14. Hopfield, J.J. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA* **81**, 3088–3092 (1984).
15. Treves, A. Graded-response neurons and information encodings in autoassociative memories. *Phys. Rev. A.* **42**, 2418–2430 (1990).
16. O'Keefe, J. & Recce, M.L. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**, 317–330 (1993).
17. Poe, G.R., Nitz, D.A., McNaughton, B.L. & Barnes, C.A. Experience-dependent phase-reversal of hippocampal neuron firing during REM sleep. *Brain Res.* **855**, 176–180 (2000).
18. Skaggs, W.E., McNaughton, B.L., Wilson, M.A. & Barnes, C.A. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**, 149–172 (1996).
19. Harris, K.D., Csicsvári, J., Hirase, H., Dragoi, G. & Buzsáki, G. Organization of cell assemblies in the hippocampus. *Nature* **424**, 552–556 (2003).
20. Skaggs, W.E. & McNaughton, B.L. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* **271**, 1870–1873 (1996).
21. Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvári, J. & Buzsáki, G. Replay and time compression of recurring spike sequences in the hippocampus. *J. Neurosci.* **19**, 9497–9507 (1999).
22. Louie, K. & Wilson, M.A. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* **29**, 145–156 (2001).
23. Paulsen, O. & Sejnowski, T.J. Natural patterns of activity and long-term synaptic plasticity. *Curr. Opin. Neurobiol.* **10**, 172–179 (2000).
24. Bi, G. & Poo, M.M. Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu. Rev. Neurosci.* **24**, 139–166 (2001).
25. Rinzel, J. & Ermentrout, B. Analysis of neural excitability and oscillations. in *Methods in Neuronal Modeling* 2nd edn. (eds. Koch, C. & Segev, I.) 251–291 (MIT Press, Cambridge, Massachusetts, 1998).
26. Brown, E., Moehlis, J. & Holmes, P. On the phase reduction and response dynamics of neural oscillator populations. *Neural Comput.* **16**, 673–715 (2004).
27. Gutkin, B.S., Ermentrout, G.B. & Reyes, A. Phase response curves determine the responses of neurons to transient inputs. *J. Neurophysiol.* **94**, 1623–1635 (2005).
28. Guevara, M.R., Glass, L. & Shrier, A. Phase locking, period-doubling bifurcations, and irregular dynamics in periodically stimulated cardiac cells. *Science* **214**, 1350–1353 (1981).
29. Ermentrout, B. & Kopell, N. Learning of phase lags in coupled neural oscillators. *Neural Comput.* **6**, 225–241 (1994).
30. Lampl, I. & Yarom, Y. Subthreshold oscillations of the membrane potential: a functional synchronizing and timing device. *J. Neurophysiol.* **70**, 2181–2186 (1993).
31. Ermentrout, G.B. & Kopell, N. Fine structure of neural spiking and synchronization in the presence of conduction delays. *Proc. Natl. Acad. Sci. USA* **95**, 1259–1264 (1998).
32. MacKay, D.J.C. Maximum entropy connections: neural networks. in *Maximum Entropy and Bayesian Methods, Laramie, 1990* (eds. Grandy, Jr, W.T. & Schick, L.H.) 237–244 (Kluwer, Dordrecht, The Netherlands, 1991).
33. Sommer, F.T. & Dayan, P. Bayesian retrieval in associative memories with storage errors. *IEEE Trans. Neural Netw.* **9**, 705–713 (1998).
34. Bi, G.Q. & Poo, M.M. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J. Neurosci.* **18**, 10464–10472 (1998).
35. Buzsáki, G. Theta oscillations in the hippocampus. *Neuron* **33**, 325–340 (2002).
36. Buzsáki, G. & Draguhn, A. Neuronal oscillations in cortical networks. *Science* **304**, 1926–1929 (2004).
37. Abbott, L.F. & Nelson, S.B. Synaptic plasticity: taming the beast. *Nat. Neurosci.* **3**, 1178–1183 (2000).
38. Ermentrout, B. Type I membranes, phase resetting curves, and synchrony. *Neural Comput.* **8**, 979–1001 (1996).
39. Engel, A., Englisch, H. & Schütte, A. Improved retrieval in neural networks with external fields. *Europhys. Lett.* **8**, 393–399 (1989).

40. Reyes, A.D. & Fetz, F.E. Effects of transient depolarizing potentials on the firing rate of cat neocortical neurons. *J. Neurophysiol.* **69**, 1673–1683 (1993).
41. Hasselmo, M.E. Neuromodulation: acetylcholine and memory consolidation. *Trends Cogn. Sci.* **3**, 351–359 (1999).
42. Froemke, R.C. & Dan, Y. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* **416**, 433–438 (2002).
43. Sjostrom, P.J., Turrigiano, G.G. & Nelson, S.B. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron* **32**, 1149–1164 (2001).
44. Lengyel, M., Szatmáry, Z. & Érdi, P. Dynamically detuned oscillations account for the coupled rate and temporal code of place cell firing. *Hippocampus* **13**, 700–714 (2003).
45. Huxter, J., Burgess, N. & O'Keefe, J. Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* **425**, 828–832 (2003).
46. Huhn, Z., Orbán, G., Érdi, P. & Lengyel, M. Theta oscillation-coupled dendritic spiking integrates inputs on a long time scale. *Hippocampus* **15**, 950–962 (2005).
47. Siapas, A.G. & Wilson, M.A. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* **21**, 1123–1128 (1998).
48. Sirota, A., Csicsvári, J., Buhl, D. & Buzsáki, G. Communication between neocortex and hippocampus during sleep in rodents. *Proc. Natl. Acad. Sci. USA* **100**, 2065–2069 (2003).
49. Csicsvári, J., Hirase, H., Czurkó, A., Mamiya, A. & Buzsáki, G. Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *J. Neurosci.* **19**, 274–287 (1999).
50. Siapas, A.G., Lubenov, E.V. & Wilson, M.A. Prefrontal phase locking to hippocampal theta oscillations. *Neuron* **46**, 141–151 (2005).