

Máté Lengyel · Zsófia Huhn · Péter Érdi

Computational theories on the function of theta oscillations

Received: 10 November 2004 / Accepted: 18 March 2005 / Published online: 17 May 2005
© Springer-Verlag 2005

Abstract Neural rhythms can be studied in terms of conditions for their generation, or in terms of their functional significance. The theta oscillation is a particularly prominent rhythm, reported to be present in many brain areas, and related to many important cognitive processes. The generating mechanisms of theta have extensively been studied and reviewed elsewhere; here we discuss ideas that have accumulated over the past decades on the computational roles it may subservise. Theories propose different aspects of theta oscillations as being relevant for their cognitive functions: limit cycle oscillations in neuronal firing rates, subthreshold membrane potential oscillations, periodic modulation of synaptic transmission and plasticity, and phase precession of hippocampal place cells. The relevant experimental data is briefly summarized in the light of these theories. Specific models proposing a function for theta in pattern recognition, memory, sequence learning and navigation are reviewed critically. Difficulties with testing and comparing alternative models are discussed, along with potentially important future research directions in the field.

1 Introduction

Neural networks, just like other dynamical systems with sufficiently many degrees of freedom, can exhibit various qualitatively different modes of asymptotic behavior. Under the biologically relevant constraint that neural activities remain bounded, network dynamics may settle to a fixed point, maintain a limit cycle oscillation or show chaotic behavior (Arbib

et al. 1998). When faced with such a repertoire of behaviors, two sorts of fundamental questions arise. Mechanistic questions ask how these different behaviors come about, i.e. how different ingredients, parameters and initial conditions of the system influence which dynamical mode is expressed. Functional questions, the subject of this review, focus on their utility, i.e. the sorts of neural computations that may be supported by one mode or the other. Traditionally, most neural network models use fixed point or line attractor dynamics to account for a large variety of computations (Amit 1989; Latham et al. 2003; Pouget et al. 2003). Less is known about the possible computational benefits of the two other modes.¹

Neural oscillations have been studied for almost a century now, but only recently has their functional significance been studied systematically (Gray 1994; Engel et al. 2003). Theta oscillations, one of the characteristic periodic electroencephalogram signals, is an exceptionally appealing subject for modeling due to the wealth of suggestive experimental data from the behavioral to the single neuron level. These data show that, on one hand, theta influences neural network activity in many ways and, on the other hand, is important for successful performance in a variety of cognitively demanding tasks. Is there a link between the two? As this type of question is not easily addressed directly by experimental techniques, it is a field ripe for theoretical investigation. Here, we review theoretical studies in which theta oscillations-specific changes in neural network dynamics were shown to contribute to meaningful computations.

In Sect. 2, we give an overview of the principal theta-related mechanisms different theories use to implement neural computations and the relevant experimental data that support these theories. In Sect. 3, we describe specific computational models that include theta as a necessary component. Our focus is on the role of theta in neural computations, and though the models discussed often have other aspects too, those are neglected here. Finally, in Sect. 4, we discuss some of the

M. Lengyel (✉)
Gatsby Computational Neuroscience Unit, University College London,
London, United Kingdom
E-mail: lmate@gatsby.ucl.ac.uk

Zs. Huhn · P. Érdi
Department of Biophysics, KFKI Research Institute for Particle
and Nuclear Physics, Budapest, Hungary

P. Érdi
Center for Complex Systems Studies, Kalamazoo College,
Kalamazoo, MI, USA

¹ Computations in the chaotic regime (or at the ‘edge of chaos’, without reaching any stable state at all) has recently become popular (Skarda and Freeman 1987; Tsuda 2001; Maass et al. 2002, but see Maei and Latham 2004).

general points and common themes that emerge from different modeling studies.

This review is limited to models dealing with the oscillatory, or phasic, effects of theta oscillations. There are also tonic effects of theta on network dynamics due to neuromodulation on a slower time scale, but these (usually cholinergic) effects are taken into account in most cases by simply setting the parameters of a non-oscillating model to specific values instead of analyzing dynamical interactions during theta oscillations (Hasselmo et al. 1995, 1996; Myers et al. 1996). Theories of this sort can be important to explore the advantage of having more than one processing modes in a neuronal structure, but tell us little about why one of these modes should endow the structure with oscillatory dynamics, and are, thus, excluded from the present discussion.

2 Experiments and theories

The theta rhythm is an appealing subject for theoretical work because it appears to be ubiquitous in the nervous system. Theta oscillation is a large amplitude local field potential oscillations around 8–12 Hz. It is present in several mammalian species, both during anesthesia or REM sleep (Jung and Kornmuller 1938; Jouvet et al. 1959) and active exploration (Green and Arduini 1954; Vanderwolf 1969). Theta is most prominent in the hippocampus but it can also be recorded in other brain regions, most notably the olfactory bulb (Adrian 1942).

Signs of its functional significance are that theta power, frequency or coherence between different brain regions has been found to covary with performance in various cognitively demanding tasks, usually involving some aspects of learning and memory (Klimesch 1999; Kahana et al. 2001). The relevance of theta in primates, let alone in the human, has been considered to be controversial (Uchida et al. 2001), but there is recent evidence of robust cortical and hippocampal theta oscillations in both monkeys (Lee et al. 2005) and humans during memory and navigation tasks (Gevins et al. 1997; Kahana et al. 1999; Tesche and Karhu 2000; Raghavachari et al. 2001; Jensen and Tesche 2002).

Since the first macroscopic descriptions of theta as a periodic signal in the local field potential or electro-encephalogram, much more fine-grained data has accumulated about the changes with which it is associated at the level of neural network dynamics. Accordingly, computational models aiming to account for the cognitive function of theta take advantage of this body of data, by making a link between some of these network-level changes and one of the well-known cognitive correlates of theta. Although these models come in many different flavors, they are most often variations on one of a few underlying theories about how theta-frequency periodicity can be exploited in a useful way for neural computations. In this section, we review these basic theoretical concepts along with the relevant experimental data that motivates and supports them. In Sect. 3, we list specific models

according to the cognitive functions in which they show improvement with the use of theta.

Note that, only few of the theories discussed below have theta oscillations as a limit cycle in the state space of their dynamical variables. They often rather have inputs or ‘parameters’ that oscillate with theta frequency, or postulate that individual neurons of the network show ‘phase precession’ which is already a corollary of hippocampal theta oscillations. These theories can thus be thought of as pertaining to brain areas receiving inputs from those actually showing theta frequency limit cycle oscillations. Nevertheless, they are still useful for elucidating possible functions that theta may play in neural computations.

Further, most of the theories could easily be generalized to other oscillations and do not necessarily need to be tied to theta or any other particular frequency. In some cases, it could be argued that only theta period length matches synaptic and cellular time constants appropriately for the proposed mechanisms to work. In other cases the reasons for which models are presented as pertaining to theta are not so much theoretical but rather practical: the experimental data they build upon were collected in connection with theta or brain regions, such as the hippocampus or the olfactory bulb, which are notoriously strong sources of theta oscillations. This does not invalidate these models, just points to the difficulty in developing theories that not only account for the functional significance of theta oscillations, but also explain why theta and not some other oscillations have those functions.

2.1 Firing rate oscillations

Historically, the computational significance of oscillatory patterns was first analyzed about 20 years ago in the context of associative memories (Baird 1986, 1990; Li and Hopfield 1989) and pattern segmentation (Wang et al. 1990). Several models of theta function investigate the computational advantage of having periodic instead of point attractors in a neural network along these lines (Sects. 3.1.3 and 3.2.2). As these models treat firing rates of neurons as state variables, they predict sustained theta-frequency oscillations in neuronal firing rates. In the hippocampus, the activity of most neurons indeed seem to be modulated by the theta rhythm at least to some degree, and different neuronal populations fire preferentially at different phases during a theta cycle (Skaggs et al. 1996; Csicsvári et al. 1999; Klausberger et al. 2003). A similar regularity was found in the entorhinal cortex where putative excitatory and inhibitory cells were also shown to differ in their preferred theta phase (Alonso and Garcia-Austt 1987; Frank et al. 2001). The activity of neurons in the medial septum is also theta-modulated, but it is less clear whether the classification based on preferred theta phase corresponds to more traditional classification schemes based on morphology or neurochemical content (King et al. 1998; Hajós et al. 2004).

2.2 Subthreshold oscillations

Another line of models breaks with the notion that the net depolarization received by a cell can be translated straight into some instantaneous firing rate, and shows that subthreshold membrane potential oscillations can interact with other sources of depolarization to produce firing patterns in which the firing times of single spikes convey meaningful information (Sects. 3.1.1, 3.1.2, 3.2.3 and 3.2.4). Theta-frequency subthreshold membrane potential oscillations were found in hippocampal pyramidal, granule and basket cells (Leung and Yim 1986; Nuñez et al. 1987; Ylinen et al. 1995; Kamondi et al. 1998), which are at least partially due to resonant properties of their membranes (Leung and Yu 1998; Pike et al. 2000). Similar subthreshold oscillations were observed in entorhinal stellate cells (Alonso and Klink 1993; Klink and Alonso 1993) and mitral cells of the olfactory bulb (Margrie and Schaefer 2003).

2.3 Periodic modulation of synaptic dynamics and plasticity

One group of neural network models of hippocampal theta oscillations contain both feed-forward and recurrent connections, and show the need to alternate between periods when one or the other synapse system dominates network dynamics or undergoes plasticity (Sects. 3.2.1, 3.3.1, 3.3.2 and 3.4.1). In line with this, both the efficacy of synaptic transmission at several hippocampal synapses and the sign and magnitude of synaptic plasticity have been shown to be modulated by theta. Current source density analysis showed periodically alternating sink-source pairs in different layers of the hippocampus where different pathways originating intrahippocampally or in the entorhinal cortex terminate predominantly (Brankack et al. 1993). The same stimulation protocol evokes postsynaptic responses of different sizes depending on the timing of the stimulus relative to the theta cycle (Wyble et al. 2000). Furthermore, a short burst of stimulation delivered at the peak of theta induces long-term potentiation of synapses, while the same burst at a trough induces no potentiation or even results in long-term depression of previously potentiated synapses (Pavlidis et al. 1988; Huerta and Lisman 1995; Hölscher et al. 1997; Orr et al. 2001).

2.4 Phase precession

The firing rate of principal cells of the hippocampus has long been known to be correlated with the position of the animal: a 'place cell' fires only when the animal is in a small portion of the environment, the 'place field' of the cell (O'Keefe and Dostrovsky 1971). More recently, it has also been demonstrated that within the place field, the timing of individual action potentials relative to the ongoing theta oscillation also varies with the position of the animal in a systematic way: place cells fire at progressively earlier phases of subsequent theta cycles, in other words they 'phase

precess', as the rat passes through their place fields (O'Keefe and Recce 1993), Fig. 1a. Several models make use of the phase precession effect showing that cells firing early or late in the theta cycle signal locations behind or ahead of the rat, respectively, and combine it with theta-modulated transmission or plasticity (as seen in Sect. 2.3) in synapses from phase precessing cells to downstream areas for decoding or learning (Sect. 3.4.2).

There is also a potentially interesting interaction between synaptic plasticity and the phase precession effect that many models exploit (Sect. 3.3.3 and 3.4.3). On one hand, while long-term potentiation and depression (LTP and LTD) were originally described as depending on the firing rates of pre- and postsynaptic neurons, more recently it was convincingly demonstrated in different preparations and brain structures, such as the hippocampus and neocortex, that the exact timing of individual pre- and postsynaptic spikes may be a fundamental factor in determining both the sign and amplitude of synaptic weight change (Paulsen and Sejnowski 2000; Bi and Poo 2001). For this to happen, pre- and postsynaptic spikes have to be synchronized in a time window of ~ 20 msec (Paulsen and Sejnowski 2000; Bi and Poo 2001). If there is just a few milliseconds jitter in spike timing within this time window, it may lead to the reversal of the order in which pre- and postsynaptic spikes are discharged. Some spike timing-dependent plasticity (STDP) rules, including those described in the cortex, are also sensitive to the order of firing: pre \rightarrow post firing results in LTP, post \rightarrow pre firing leads to LTD (Markram et al. 1997; Bi and Poo 1998). Therefore, cortical plasticity seems to be extremely sensitive to the precise timing of neuronal firings.

Phase precession, on the other hand, has a profound effect on the cross-correlation of spike trains fired by place cells having neighboring place fields (Skaggs et al. 1996). A place cell fires through several theta cycles while the animal passes through its place field and in each subsequent cycle it fires at an earlier phase as dictated by the phase precession effect (O'Keefe and Recce 1993; Skaggs et al. 1996). If two cells have their place fields after each other along the trajectory of the animal, they will fire together during several subsequent theta cycles, and the cell having the second place field will always fire at a later phase, because it only starts precessing later (Fig. 1B). Consistent with this, such pairs of place cells show significant peaks in their spike train cross-correlation histograms, the first peak appearing within a theta period and then reappearing with theta frequency (Skaggs et al. 1996; Dragoi et al. 2003; Harris et al. 2003). Therefore, by compressing temporal sequences and preserving the sequence in which place fields were visited in the compressed sequence of neuronal firings, phase precession provides both the time window and consistency of cell firing ideal for STDP to be effective. The experience-dependent expansion and backward shift of place fields (Mehta et al. 1997) have been interpreted as a result of precisely this effect: phase precession with STDP consistently strengthens the connections between place cells having subsequent place fields along the trajectory of the animal (Mehta et al. 2002).

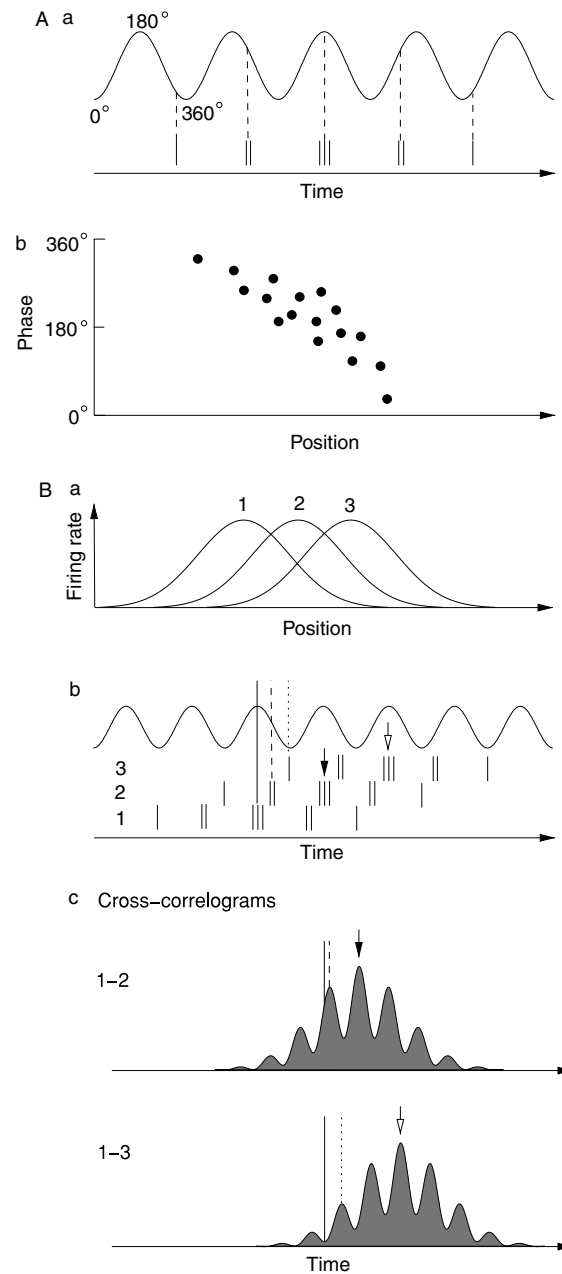


Fig. 1 The phase precession effect in hippocampal place cells. **A** Phase precession in one cell. **a** The firings of a place cell (*bottom spike train*) occur at progressively earlier phases relative to the ongoing theta field potential oscillation (*top sinusoid trace*) as the rat traverses the place field of the cell. Note that the number of spikes fired per cycle also varies (see also subfigure *Ba*). **b** Scattergram shows the phase of theta oscillation (*y-axis*) and the position of the animal as it runs from left to right (*x-axis*) at the time of individual spikes (*dots*) collected through several traversals of the place field. Although the phase of spikes is also correlated with the time since entry to the place field, correlation is stronger with the position of the animal (O'Keefe and Recce 1993). The extent of phase precession can reach 360° , i.e., a full theta cycle. **B** Phase precession in cells with overlapping place fields. **a** The number of spikes fired per cycle (*y-axis*) varies according to a unimodal function of position (*x-axis*): firing rate is low near the beginning and end, and peaks in the center of the place field. The three place cells (*1, 2, 3*) shown have overlapping place fields, the animal runs from left to right. **b** Spike trains of the three cells (*bottom rows*) are shifted relative to each other, but phase precession is similar in all cases: its starting and ending phase is near 0° , and it is centered around 180° . Time differences between the *solid line* and the *filled* or *open arrow* show the overall time shift of the spike trains of cell 2 and 3 relative to the spike train of cell 1. Time differences between the *solid* and the *dashed* or *dotted line* show consistent phase lags at which cell 2 or 3 fires behind cell 1 within each theta cycle due to the phase precession effect. **c** Cross-correlograms of spike trains after Skaggs et al. 1996. Cross-correlograms are theta-modulated due to phase precession and have a unimodal envelope due to the unimodal firing rate tuning curve of place cells. The distance of overlapping place fields (*1-2*, and *1-3*) is reflected both in the shift of the center of the envelope from zero (*filled* or *open arrow* from *solid line*) and, on a compressed time scale, in the latency of the first peak in the correlogram (*dashed* or *dotted line* from *solid line*). This and the following figures are intended for illustration only and were not constructed from experimentally recorded or simulated data

3 Models

3.1 Pattern recognition

Pattern recognition is a classical challenge for neural networks (Rosenblatt 1958; Duda et al. 2001). In its canonical form, a network has to classify a highly distributed pattern of input activity (such as pixel values in an image) into one of a small number of possible categories (such as houses vs. faces). In the case of regression, the network is also allowed to give answers like ‘70% house, 30% face’. Traditionally, inputs to a neural network are thought of as distributed patterns of injected current: each neuron is depolarized to a different level, and the output of the network is given in terms of the firing rates of some of its neurons.

The first two models in this section explore the possible benefits of representing outputs in terms of relative spike times (spike phases) rather than firing rates, when there is also a synchronous subthreshold oscillating input to all cells of the network. These models use feed-forward networks. The third model investigates recurrently coupled networks using firing rates codes, and points out that limit cycles can alleviate the trade-off between effectively amplifying a selected subset of inputs and avoiding the spontaneous amplification of noise into ‘hallucinations’ when there is no meaningful signal in the input.

3.1.1 Concentration invariant pattern recognition

It was first proposed by Hopfield (1995) purely on theoretical grounds that subthreshold membrane potential oscillations may endow neurons with a temporal code that can effectively support a specific type of computation apparently important in olfactory processing: concentration invariant odor recognition. In the model, cells receive a subthreshold oscillatory input plus a depolarizing current (Fig. 2a). Depending on the level of depolarization a cell will fire earlier or later during the theta cycle – as long as the depolarization remains in the regime in which the cell fires exactly one spike per cycle. This way, the timing of action potentials relative to the theta oscillation will depend quasi-logarithmically on the amount of depolarizing current (Fig. 2c). If a given odor excites several such cells with different depolarization levels proportional to the concentration of the odorant, then these cells will fire at different times during the subthreshold oscillation, but due to the logarithmic mapping from depolarization to spike timing, the timing of spikes fired by different cells relative to each other will be preserved, independent of odorant concentration. An appropriate decoder, sensitive for a given pattern of relative presynaptic spiking times (e.g., by being connected to presynaptic cells through delay lines of different lengths) would then be able to recognize odors in a concentration invariant way. Brody and Hopfield (2003) further elaborated on this idea, showing that such a network is also capable of odor segmentation (recognizing individual odors in a mixture) and odor recognition in the presence of strong distractor odors. Hopfield (2004) also showed how essentially the same

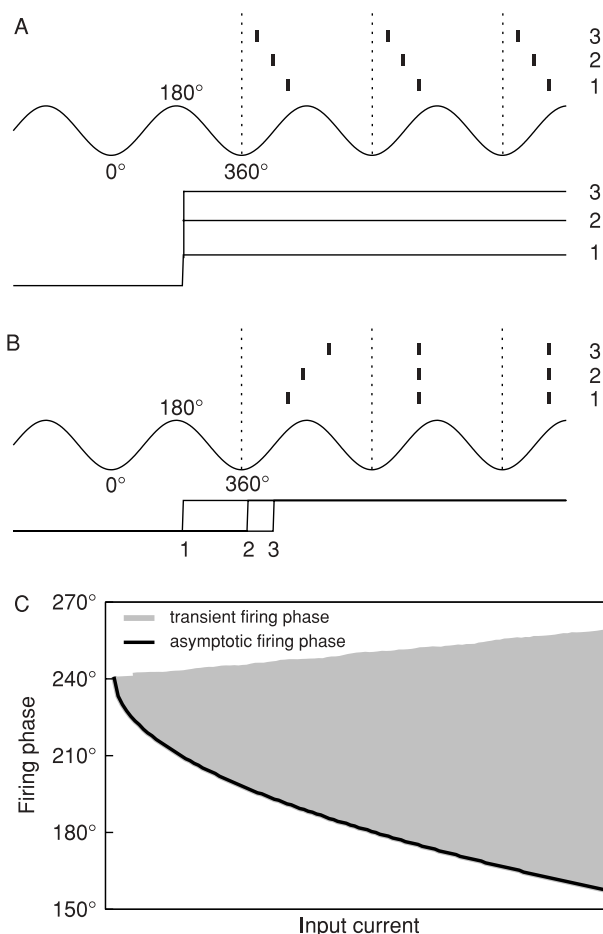


Fig. 2 Encoding depolarization level in firing phase. **a** The timing of spikes (top spike trains) within the theta cycle (middle sinusoid trace) depends on the level of depolarizing current (bottom step functions) received by the cell (Hopfield 1995). Spike trains labeled with 1, 2, 3 are triggered by increasing input step currents 1, 2, 3, respectively. Relatively low depolarization (1) causes the cell to fire near the peak of theta, increasing input current (2, 3) shifts firing times towards the negative peak of theta (dotted line). **b** When depolarization changes, firing times in the first cycle after the change (transient firing phase) may differ from firing times in subsequent cycles (asymptotic firing phase) (Lengyel and Érdi 2004). Spike trains labeled with 1, 2, 3 are triggered by time-shifted input step currents 1, 2, 3, respectively. Timing of the first spike depends on the exact phase of current switch-on: for some switch-on times (1) it is identical to the asymptotic firing phase, for others (2, 3) the first spike occurs later than the asymptotic firing phase. **c** Asymptotic firing phase depends logarithmically on the depolarizing current (black solid line). Firing during transients can occur in a broader phase range (gray shaded area); the size of this range increases with the level of current (modified from Lengyel and Érdi 2004)

network could be used to recognize brief dynamical patterns in time-varying stimuli, such as syllables in speech.

The biological plausibility of this mechanism has recently been confirmed by Margrie and Schaefer (2003) in *in vivo* recordings in the mouse olfactory bulb. The existence of a strong theta modulation of mitral cells was demonstrated that regulated single cell firing, just as predicted by theoretical studies. Furthermore, the feasibility of spike latency coding

based on data from the measurements was assessed, and in a direct comparison with alternative coding schemes, utilizing only binary glomerular activation patterns or the graded rate code of individual mitral cells, it was proved to yield significant computational power.

3.1.2 Detecting input transients

Lengyel and Érdi (2004) pointed out that the strict logarithmic dependence of firing time on depolarization level assumed in earlier works is only true asymptotically, i.e., after the cell has been depolarized at the same level for a sufficiently long time (Fig. 2b). When the level of depolarization changes, firing times in the first theta cycle, but not in later cycles, following such a transient can significantly deviate from the asymptotic firing time and are not captured by the logarithmic relationship. In particular, while asymptotic firing phases are restricted to the depolarizing half of the theta cycle, transient-elicited spikes may appear at a broader range of firing phases (Fig. 2c). Thus, a mechanism that is sensitive to firing phases may be able to discriminate firing during transients from firing during steady-state inputs. This provides a ‘spike-jitter’ code for transients in the input. This code can be read out by postsynaptic cells if transmission from the spike timing-coding neurons is only allowed during the hyperpolarizing half of the theta cycle. Note, that in this case theta is necessary not only for the encoder but also for the decoder part of the network.

The authors also proposed that this scenario may be relevant for the entorhino-hippocampal system, entorhinal and hippocampal principal cells playing the roles of the encoder and the decoder, respectively. This is because both cell populations are modulated by a coherent theta oscillation (Skaggs et al. 1996; Frank et al. 2001), and synaptic transmission between the two populations has been shown to be effectively regulated by perforant path-specific interneurons of the hippocampus (Freund and Buzsáki 1996), which in turn are also appropriately theta-modulated (Csicsvári et al. 1999; Klausberger et al. 2003).

3.1.3 Selective amplification while avoiding spurious attractors

Selective amplification can be considered to be a special case of pattern recognition, when some components of the input (usually considered to be the signal) are transmitted to the output in a highly amplified form, while others (comprising noise) get almost completely suppressed. This situation is further complicated when the same motif in the input (such as a vertical line in a visual scene) should be amplified in some cases (e.g., a stand-alone line indicating the contour of an edge of an object) and should be suppressed in other cases (e.g., a line among other lines of the same orientation composing the texture of an object). To achieve this, neural networks need intricate excitatory-inhibitory interactions. Unfortunately, such excitatory-inhibitory networks are prone to hallucinations (Ermentrout and Cowan 1979): strongly

inhomogeneous output patterns form spontaneously due to spurious attractors in the dynamics of the system in response to near-homogeneous input (that is supposed to convey no meaningful information and therefore should lead to overall suppressed output). Thus there appears to be a trade-off between selective amplification and avoiding hallucinations.

Li and Dayan (1999) systematically analyzed the dynamical behavior of excitatory-inhibitory systems in which excitatory and inhibitory units come in pairs, and these pairs are interconnected through symmetric long-range excitatory connections targeting both excitatory and inhibitory units. In such systems, the amplification-hallucination trade-off holds as long as the membrane time constant of inhibitory cells is negligible compared to that of excitatory cells thus making the effective synaptic weight matrix of the system symmetric. A completely different behavior emerges, however, if these time constants are comparable and thus the effective synaptic weight matrix becomes asymmetric. In this case, the network shows limit cycle oscillations, the response of the system being simply the amplitude of these oscillations, and thus avoids sliding into unwanted fixed point attractors. This network not only is saved from spurious firing, but the degree of selective amplification it can achieve can also be orders of magnitude larger than in an otherwise equivalent non-oscillating system. Although the frequency of limit cycle oscillations depends on model parameters, there seems to be an interesting relationship between the amplitude and frequency of these oscillations in the model: amplification of the input results in lower frequencies than suppression.

As the authors do not commit themselves to any particular frequency band, this model is not specifically a model of theta oscillations. In fact, gamma oscillations could be the most likely candidate to be the corresponding biological rhythm due to its prevalence in the visual system, to which this model originally was applied, and its proposed role in signalling ‘recognition’ of parts of a stimulus belonging together. However, the computational role that limit cycle oscillations in this model subserve appears to be rather general, that may be necessary to implement in several brain areas, and the structural ingredients that these oscillations require also appear to be generally present in the nervous system. Therefore, it may be the case that theta oscillations in other brain areas have a similar function.

3.2 Memory

Neural networks models of memory are fundamentally related to those of pattern recognition. The network is first trained to store some patterns by changing the synaptic efficacies in the connections between its constituent neurons. Then in the test phase the network is given a genuine classification task: it is presented an input pattern and has to recall that pattern in the training set which is the most similar to its input. However, while pattern recognition models described above assumed that the synaptic weights were already set in the system appropriately, memory models are concerned with

how storing patterns in the training phase influences recall in the test phase.

The first model in this section addresses the problem of separating storage and recall, and how theta can help with it. The second group of models explore the possible benefits of storing memories as limit cycles rather than fixed points in state space, when state variables are firing rates (or membrane potentials). Third, two models are discussed which consider the case when memories are fixed points, but state variables are firing phases rather than firing rates. Finally, an oscillatory short-term memory model is discussed, which is fundamentally different from all the above models, as it is modeling a form of memory which is stored in reverberating neural activity rather than synaptic weights.

3.2.1 Preventing encoding and allowing unlearning during recall

Models of memory are often faced with the need for separate training and testing phases, when the encoding and recall of memory traces occur, respectively. Recall during encoding, when there is not much yet that could be sensibly recalled, can contaminate memory by storing nonsense activity patterns (see also below at sequence learning), and encoding during recall may lead to overlearning by repeatedly storing patterns already learned and thus displacing other stored patterns.

The problem of encoding during recall becomes especially pronounced in a reversal learning paradigm, in which the task is first to store an association between two items (going to a location and finding reward there) and then to 'forget' this association (because there is no reward anymore at the original location) and to learn a new association between other two items (going to another location and finding reward there). Performance in this task was shown to be hippocampus-dependent, motivating Hasselmo et al. (2002) to develop a model in which input to hippocampal areas CA1 and CA3 from entorhinal cortex (EC), synaptic transmission in Schaffer collaterals from CA3 to CA1, somatic membrane potential of CA1 cells, and LTP in Schaffer collaterals were all theta-modulated. The phase relationship between these periodic modulatory effects optimizing performance in terms of expressing the new association instead of the old one after reversal was analytically derived and found to be consistent with experimentally found values. Specifically, maximal EC input had to be in phase with maximal LTP, and both needed to be in antiphase with maximal transmission in Schaffer collaterals and somatic depolarization in CA1 cells. The alternating activation of these two subsystems lead to separate encoding and retrieval phases. A key element in the simulations was that LTP was turned to LTD at negative phases of its modulation, i.e., during the retrieval phase, so that associations not presented during encoding phases were quickly unlearned by an anti-Hebbian learning rule.

The model matches with experimental data obtained both on electrophysiological and behavioral levels, yet, it can be criticized from a strictly computational point of view. On

one hand, it is unclear how the system would benefit from the same mechanisms (LTD during retrieval) in any one of the numerous other hippocampus-dependent learning paradigms not involving reversal learning. On the other hand, it is easy to imagine that the same task of reversal could be learned by a simpler model without oscillatory dynamics (e.g., two presynaptic cells synapsing on a single postsynaptic cell with competing synapses).

Fast alternation between encoding and retrieval as opposed to discrete training and testing stages, however, could be important and advantageous for a neural network, and it would be exciting to see in computational models how such a mechanism could be exploited in a sensible way. Perhaps, the alternating Hebbian and anti-Hebbian phases of contrastive divergence used for statistical learning in the restricted Boltzmann machine (Hinton 2002; Káli and Dayan 2004) could be implemented by the theta rhythm.

3.2.2 Storing oscillatory patterns

Traditional autoassociative networks store memories that are represented as distributed patterns of stable firing rates. If the state variables of the neural network are the firing rates of units, stored memories will be fixed points in state space. The following models explore the possibility when memories are represented as oscillatory patterns and are therefore limit cycles in state space.

Borisuyuk and Hoppensteadt (1999) propose that theta-modulated input from the septum and from EC creates specific spatio-temporal activity patterns in the hippocampus, where single cells, or interconnected excitatory-inhibitory populations (Wilson and Cowan 1972), act as intrinsic oscillators at the gamma frequency. Through complicated interactions of the two external and the intrinsic oscillations, various interference patterns emerge, depending on the exact connectivity patterns, phase delays in afferent fibers, and other parameters. These patterns were shown to be storable in the hippocampal network by a Hebbian learning rule sensitive for the synchrony of pre- and postsynaptic cells (Borisuyuk and Hoppensteadt 1998). If appropriately trained, these patterns can be recalled as a sequence, although there must be considerable overlap between them during training. However, it is not clear what the network benefits from the memories being stored as interference patterns and not as arbitrarily defined patterns over cells. Indeed, in the first model dealing with the storage of patterns (Borisuyuk and Hoppensteadt 1998), input from EC to be learned is provided as a series of static binary activity vectors and not through interference with septally generated oscillations. In a more recent version of the model (Borisuyuk et al. 2000) both inputs are oscillatory, encoding and recall works through resonance and learning tunes the natural frequencies of the intrinsic oscillators. It is shown that this network may signal the novelty of a stimulus by the number of recruited resonant units in response to the stimulus. Yet, neither the biological plausibility of the network, nor the direct computational benefits of using an intricate oscillatory-resonant mechanism for novelty detection in

comparison with a simpler rate coding scheme (Bogacz et al. 2001; Legenstein et al. 2003) seem to be firmly established.

Scarpetta et al. (2002) used a similar network architecture (that is also similar to the model of Li and Dayan (1999)): coupled excitatory-inhibitory populations formed pairs receiving oscillatory input, and excitatory cells of each pair also connected to both excitatory and inhibitory cells of other pairs. Both of these types of long-range connections were allowed to change with Hebbian STDP, which in turn modified the resonant properties of the network so that it became specifically resonant to those patterns that had previously been stored. This could be achieved only if the shape of the STDP curve satisfied some constraints, and realistic STDP shapes were indeed found to be 'optimal' in this sense. At recall, this network could perform pattern completion on oscillatory patterns in the same way traditional autoassociative networks complete distributed patterns of firing rates. Moreover, compared to traditional networks it had increased input specificity, because both the frequency of the input oscillation, which was the same for all units in the network for a given pattern, and the phase of the oscillation, which could be different for each unit, had to match those of a stored pattern. This potentially yields superior performance over traditional networks, because increased input specificity reduces interference between different stored patterns, which is known to ultimately limit the capacity of autoassociative networks (Hopfield 1982). However, the capacity of the network was not systematically studied, and in fact the analytical derivations only guarantee proper functioning of the network in the limit where only very few patterns are stored. Nishikawa et al. (2004) analyzed the capacity of oscillatory autoassociative memories and found that it was close to that of the classical rate-based Hopfield network.

Another interesting finding by Scarpetta et al. (2002) was that changing the nonlinearity in the frequency-current transfer function of cells could switch the network between two qualitatively different modes of operation. In one mode, the network acts as a classifier on its inputs: it can only relax to oscillations that correspond to previously stored patterns. In the other mode, it performs regression: it can also relax to oscillations which are mixtures of previously stored patterns, thus effectively interpolating between them. Although, just as in Li and Dayan (1999), the oscillations are not necessarily in the theta band, the authors point to those two areas as possible biological substrates of their model where theta is the most prominent: the olfactory bulb and the hippocampus. They propose that different neuronal nonlinearities in the two areas could lead to classification in the olfactory bulb and regression in the hippocampus, or that neuromodulators could switch between modes within the same area by tuning the transfer function of neurons.

3.2.3 Storing distributed phase-coded memories

The work of Wills (2004) (also MacKay and Wills, unpublished data) explores the possibility that neural networks may store distributed patterns of firing times relative to the theta

oscillation, i.e. firing phases, rather than distributed patterns of firing rates. By further elaborating on the spike timing-coded network of Brody and Hopfield (2003), they propose that decoding neurons receiving the same subthreshold oscillatory input as encoding neurons will signal recognition of an input pattern by the means of the same temporal code that is used by encoding neurons to signal the presence of components of a pattern. Encoding and decoding neurons thus using the same code need not be distinguished and separated into two layers of a feed-forward network anymore, but can be treated as a homogeneous population of recurrently coupled neurons. This turns pattern recognition of the feed-forward network into pattern completion: a neuron recognizing a fragment of an input pattern emits a spike, which together with spikes from other neurons is recognized by another neuron also emitting a spike, and so forth, thus giving rise to a complete pattern of firing times. This mechanism seems to be able to perform operations that are traditionally considered for rate-based networks to be hard to cope with. Analog patterns, in which unit activations need not be binary, can be recalled with great precision, multiple patterns can be recalled simultaneously, and operating with line attractors becomes a straightforward extension of the theory.

However, for the model to work, memories need to be encoded in the connection delays between cells, which seems to be at odds with the conventional view that memories are encoded in synaptic efficacies. Although the authors offer a scheme in which dendritic spines and main branches could support this computation, it is unclear whether there is really the biological machinery required by the theory. Alternatively, Lengyel and Dayan (2005) show that conventional Hebbian STDP storing memories in synaptic efficacies can be used to store distributed phase-coded memories if matched by appropriate neuronal interactions during recall. The theory predicts a specific phase response curve² that is optimal for recalling memories stored by a given form of STDP. This provides a consistent framework for studying the storage and recall of phase-coded memories.

3.2.4 Sustaining ordered firing for short-term memory

The previous models were concerned with long-term memory, a form of memory that is stored in synaptic weights and can be recalled at a time that can be (in theory) arbitrarily later than storage. Short-term memory is usually thought to be maintained by reverberating neural activity, and is shown to be sensitive to distractors which poses serious limits on the maximal time gap between storage and recall. (Its prototypical example is recalling a telephone number just looked up from the phone book.)

The model of Lisman and Idiart (1995) claims to account for the well-documented 7 ± 2 capacity of human short term

² Phase response curves describe the amount by which firing in a presynaptic cell delays or advances firing in an otherwise regularly firing postsynaptic neuron as a function of the time difference between the presynaptic spike and the previous postsynaptic spike (Rinzel and Ermentrout 1998).

memory. A feed-forward network of cells with recurrent inhibition is modeled. Each neuron receives the same theta-frequency subthreshold depolarization and also generates an afterdepolarizing potential (ADP) after emitting a spike. If the time course of the ADP matches that of the theta cycle then the depolarizing phase of these two inputs will coincide and the cell will fire in the next theta cycle, but not before that, without sustaining the input that elicited the first spike. Firing of a cell delays the firing of all other cells that did not fire together with it in perfect synchrony through lateral inhibition, therefore only one highly synchronized group of cells may fire in each time window corresponding to the time course of fast inhibition. This results in faster (approximately gamma-frequency) oscillations nested within each theta cycle, one cell group firing per gamma subcycle. The two oscillations together lead to repetition of cell firings in the same sequence in consecutive theta cycles. The capacity of this short-term memory buffer is determined by the ratio of the periods of nested theta and gamma oscillations. This model could also account for psychophysical data (e.g., response delays) in the Sternberg task (Jensen and Lisman 1998). It was later also combined with a long-term memory network to make it more robust to noise, and to allow overlap between individual memories (Jensen and Lisman 1996a,b,c, 1998; Jensen et al. 1996). Furthermore, this network was supposed to provide several hippocampal sequence learning (Sect. 3.3.1) or navigational (Sect. 3.4.1) models with appropriately pre-processed input (Jensen and Lisman 2005).

Theta oscillations were also implicated in transitive inference during paired associate tasks (Wallenstein et al. 1998). In this paradigm, the animal is trained on pairs of associations like 'if A then B' and 'if B then C'. Transitive inference occurs when the rat chooses 'C' when presented with 'A', thus expressing associations it has never encountered before. The mechanisms by which theta may aid performance in this task are equivalent to those used for sequence learning and are discussed below.

3.3 Sequence learning

The paradigm of sequence learning represents an intermediate stage between episodic memory and navigation. In its simplest form, the network is first trained on a sequence of activity patterns, by simply giving subsequent patterns as its input in subsequent time steps, and then it is tested to recall these patterns, preferably in the order in which they were presented, by priming it with only the starting element (or few elements) of the sequence as a recall cue. This can be interpreted as a serial recall of subsequent episodes, making transitive inference on paired associates, or recalling locations subsequently visited along a trajectory through space.

There is considerable evidence that the hippocampus stores sequences. The theta phase precession of place cells itself has been interpreted as resulting from the read-out of previously stored sequences (Skaggs et al. 1996; Tsodyks et al. 1996; Jensen and Lisman 1996a,b,c, 1998; Wallenstein and Hasselmo 1997), cf. with models discussed below, interpreting

phase precession as necessary for the *storage* of memory sequences. Sequences of neural activities recorded in the hippocampus of the awake animal are also replayed during subsequent sleep sessions (Skaggs and McNaughton 1996; Nádasdy et al. 1999; Louie and Wilson 2001; Lee and Wilson 2002), giving further support to sequence learning theories.

The first two groups of models reviewed here propose a direct way in which theta could interact with network dynamics to improve the performance of the network. The third type of model uses theta indirectly by exploiting the phase precession effect in hippocampal place cell firing.

3.3.1 Preventing interference

The model of Wallenstein and Hasselmo (1997) proposes that the level of γ -aminobutyric acid (GABA) changes rhythmically during theta oscillations, and focuses on the effect of periodic GABA_B receptor activation-induced suppression of intrinsic connections of region CA3 (recurrent collaterals between pyramidal cells, and connections between interneurons and pyramidal cells). This periodic modulation of intrinsic synaptic strengths results in feed-forward connections from EC dominating the early phases of the theta cycle, while intrahippocampal interactions driving the network during late phases.

The network benefits from this periodic change during both training and testing. When the network is trained on a new sequence, if intrinsic connections were always active, then activation patterns elicited by each element in the sequence would interfere with activation patterns emerging due to the processing of previous elements in the sequence by intrinsic connections. Because synapses in intrinsic connections are strengthened by a Hebbian STDP learning rule, i.e., every time when postsynaptic cell activation follows presynaptic activation within a limited time window, the model learns this mixture of externally imposed and internally generated patterns. The 'GABA up' phase of theta oscillations prevents recall during encoding by providing time windows in which intrinsic connections do not influence network dynamics significantly so that input patterns presented during these time windows are not corrupted and thus can be successfully stored. Conversely, for associations to be made between cells that are not excited directly by the input and thus for 'context fields' to emerge, a key element for the functioning of this model, it is also important that during training there is a part of the theta cycle, the 'GABA down' phase, in which internal propagation of activity does not have to compete with external input (Wallenstein et al. 1998).

During testing, the network is presented with a pattern from the stored sequence, and if this cue is active just when GABAergic modulation starts to decline, 'GABA down', then activation spreading along intrinsic connections leads to the recall of subsequent elements in the sequence. If GABAergic suppression does not decrease sufficiently in the second half of the cycle, no recall of the sequence occurs. If, on the other hand, GABAergic suppression of intrinsic connections is not strong enough at the beginning of the

cycle, ‘GABA up’ phase, then external input must compete with internal dynamics, as during training described above, and cannot prime the network effectively. Consequently, sequence retrieval is impaired without a reliable recall cue.

A different type of interference is prevented by theta modulation during encoding in the sequence learning model of Jensen and Lisman (1996a,b,c). Here, a whole sequence is encoded within a single theta cycle by clamping the activity of neurons to subsequent inputs representing elements of the sequence. This encoding is repeated in several consecutive theta cycles. Association between neurons coding subsequent patterns are stored again by a Hebbian STDP learning rule using slow NMDA channels with a time constant matching a theta period. If input to the network is not interleaved with silent intervals at theta frequency, then NMDA channel activation caused by the last element of the sequence will not yet have declined sufficiently, and when the first element is presented again it gets associated with the last element. Theta prevents this breaking of directionality of the sequence so that real sequences and not only loops can be stored.

3.3.2 Providing simulated annealing dynamics

In their next model, Sohal and Hasselmo (1998a,b) addressed the problem of forking sequences, the so-called ‘sequence disambiguation problem’. If two or more sequences share the same initial pattern, or sub-sequence, then it is ambiguous which branch should be recalled when cued with the common initial element. Here the possibility was explored that a small subthreshold bias provided by external input to the final element of one of the sequences is able to direct recall to that branch during the testing phase in a model framework similar to that used in their previous work. (An analogy would be to think of a rat at a branching point trying to find out which way to take to get to the food source.) Neurons belonging to different branches reciprocally inhibited each other through common feed-back interneurons, while neurons belonging to the same branch excited each other through recurrent connections, thus resulting in a competition between branches for being activated. GABAergic modulation regulates this competition by tuning the efficacy of all (excitatory and inhibitory) intrinsic connections. During the ‘GABA up’ state, cells representing both branches get activated due to feed-forward and weak recurrent excitation, but as GABA levels decline competition starts that is essentially decided by the small bias through a ‘winner-take-all’ mechanism by the end of the cycle. This process was shown to be mathematically analogous to simulated annealing, a process known to be suitable for finding states with globally minimal energy when many local energy minima exist (Sohal and Hasselmo 1998b). Gradual introduction of competition between stored sequences as provided by theta regulated GABA_B receptor activity levels was also found to be important when one of multiple non-overlapping sequences had to be recalled in the previous model (Wallenstein and Hasselmo 1997).

The same model was demonstrated to successfully reproduce transitive inference in a paired associate learning paradigm (see above) that was observed to be hippocampus-dependent (Sohal and Hasselmo 1998a; Wallenstein et al. 1998).

3.3.3 Facilitating one-shot sequence storage

Phase precession provides ideal conditions for forming association through a STDP learning rule between cells representing consecutive episodes or subsequent locations (Fig. 1b). Yamaguchi (2003) showed in a model of sequence learning that the repeated activation of such cell pairs with consistent phase lags due to phase precession led to an asymmetric synaptic weight matrix in the recurrent connections of CA3 even after a single presentation of a sequence. This model was further analyzed (Sato and Yamaguchi 2003) to show that even cells representing elements of the sequence which are separated by a delay several times larger than the time scale of the learning rule get associated so that the sequence can be recalled successfully. When compared with a rate model (Sato and Yamaguchi 2003), the model based on phase precession yielded a higher degree of robustness to noise and less sensitivity to the overlap and separation between subsequent patterns in the sequence. However, even the phase precession model required some degree of overlap between subsequent patterns, a restriction that may hinder it in serving as a store for episodic memories. Moreover, the models are built on the assumption of phase precession being generated primarily in EC cells (and hippocampal cells merely ‘inheriting’ this precession through the perforant path) which has no experimental support to date.

Note that models without assuming any oscillatory effects were also developed for sequence learning. These models usually require larger overlaps in the set of activated neurons between subsequent patterns both for successful sequence recall and for disambiguation (Levy 1996; August and Levy 1999), or alternatively the time constant of synaptic modification have to be matched to the duration for which each element in the sequence is presented (Mitman et al. 2003). These models demonstrated successful disambiguation of looping, intersecting, or overlapping rather than forking sequences, which is based on different past activity sequences rather than biasing inputs (Levy 1996).

3.4 Navigation

Perhaps the most prominent behavioral correlate of hippocampal theta oscillations is when the rat engages in active exploration of the environment. This seems to be in accord with the cognitive map theory of hippocampal function (O’Keefe and Nadel 1978) which sees hippocampus primarily as an organ supporting navigation (at least in rodents). The following models explore how theta could aid navigation of the animal.

Here again, some models use theta-controlled network dynamics while others use the phase precession effect to demonstrate a functional role for theta in neural computation.

3.4.1 Coordinating activity propagation in EC and CA3

The model of Hasselmo et al. (2002) and Koene et al. (2003) addresses the problem of finding the shortest path to a goal in an environment where multiple goals exist and each of them can be approached through potentially multiple paths. Both EC (layer III) and CA3 stores paths in the environment (by a mechanism essentially equivalent to a sequence learning mechanism), with sequences being stored in their original order in CA and in a reversed order in EC. During navigation, the process of taking the next step begins with EC activity propagating from the location of the goal (or goals) backwards along already stored paths. When the CA3 cell representing the current location of the animal gets excited by this input a limited forward spread starts in CA3, thus recalling the locations that can be directly approached from the current location along known paths. These two activity ‘waves’ have to synchronously excite a CA1 cell to fire it. The authors demonstrate by mathematical analysis that this mechanism ensures that the CA1 cell firing will represent the next location towards the nearest goal along the shortest path. Theta is essential in two ways for this process. First, it should prevent recall during encoding, just as proposed in sequence learning models, otherwise the activity propagation in CA3 during encoding is also stored, and the location recalled during navigation will contain this propagation term, thus representing a location that is not adjacent to the present one and leading to the so-called ‘skip ahead’ problem. Second, activity propagation in EC and CA3 during recall must be appropriately timed. Theta oscillations are proposed to provide this timing, but the precise mechanism for that is not outlined.

3.4.2 Separating signals of upcoming and past locations

A consequence of phase precession not emphasized until now is that cells firing late in the theta cycle tend to have place fields *ahead* of the rat, because the animal is just entering their place fields, while those firing early in the cycle have their place fields *behind* the rat. This effect was exploited in a neural network simulation by Burgess et al. (1994) for guiding the navigation of the rat to goals located in a two-dimensional environment (Fig. 3). In the model, each subicular cell has a place field that is essentially the sum of the place fields of a set of hippocampal place cells, resulting in large subicular place fields, consistent with experimental data. Moreover, subicular cells are hypothesized to show phase precession as ‘inherited’ from hippocampal cells through feed-forward connections. A population of ‘goal cells’ is postulated one synapse downstream from the subiculum. A goal cell fires when the rat encounters a goal and, due to postulated input from head-direction cells, the rat heads into a direction (North, West,

etc.) specific for the given goal cell. Synapses between subicular goal cells are strengthened by a Hebbian learning rule requiring synchronous pre- and postsynaptic activity. During navigation, information flows from hippocampal place cells, through subicular cells, to goal cells. The population vector represented by goal cells should indicate the relative direction (and inverse distance) of the rat from the goal. However, if the phase code of subicular cells is not taken into account at learning, then goal cells will have equal input from subicular cells that had place fields in front of and behind the rat when it encountered the goal. Consequently, their firing field will be on average centered on the goal and will not convey any information about the relative direction of the rat from the goal. If change in synaptic weights is only allowed to occur in the second half of the theta cycle, then goal cells will have a bias from subicular cells having place fields in front of the animal when it encountered the goal, resulting in a place field shifted in the preferred direction of the goal cell relative to the position of the goal. Thus, during navigation the activity of goal cells will predict the direction of the goal.

The same principle was studied in a more general framework by Jensen (2001). Here, a decoder layer of spiking neurons was supposed to receive input from phase coded hippocampal place cell firing. The decoder was modulated by the theta oscillation (as a simple injected current) lagging behind hippocampal theta with an adjustable phase lag. It was shown that depending on the phase lag between encoder (hippocampus) and decoder regions, activity in the latter signals upcoming or current locations.

3.4.3 Trajectory learning

The model of Burgess et al. (1994) was further developed by Trullier and Meyer (2000). In this model, connections between (CA3) place cells were also considered, and supposed to undergo Hebbian plasticity during exploration. However, these recurrent connections were ‘labeled’ with respect to the current heading direction of the animal, so that a synapse between two cells was strengthened if the presynaptic cell fired shortly before the postsynaptic cell *and* the animal was heading in the direction corresponding to the ‘label’ of the synapse. It seems to be hard to justify the biological plausibility of such an assumption, but it proved to be efficient in guiding navigation. When the rat encounters a goal, goal cells get activated as in the model of Burgess et al. (1994), but, instead of simply activating place cells representing the current location this model starts propagating activity in the CA3 network, similarly as in the model of Hasselmo et al. (2002) discussed above, along synapses labeled with the current heading direction of the animal. The rat is supposed to probe several directions at a goal, thus recruiting place cells at different directions from the goal. All place cells thus activated are associated with the currently active goal cell, effectively enlarging the ‘prediction radius’ of the animal within which it can tell the direction of a goal relative to its current location during future navigation. Note that phase precession in this model was not necessary for

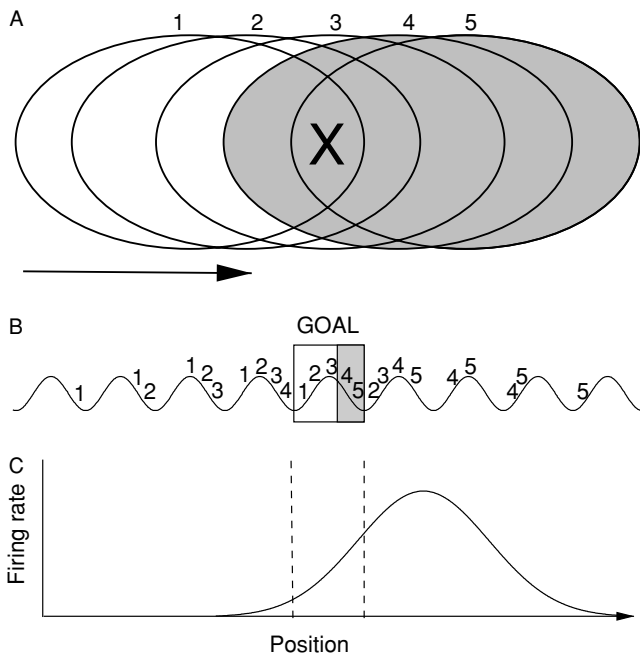


Fig. 3 Learning to predict goal locations using phase precession (Burgess et al. 1994). **a** Position of overlapping subicular place fields (ellipses) belonging to subicular cells 1, 2, 3, 4, 5. A goal cell is active when the rat is at the goal site (X) and heads into the cell's preferred direction. Here, the rat is running from west to east (arrow), so only 'East' goal cells are active near the goal. **b** Subicular place cells are hypothesized to phase precess (numbers show corresponding cell firing) relative to theta oscillation (sinusoid trace) similar to hippocampal place cells: a cell firing late in the theta cycle has most of its place field ahead of the rat's current location. 'East' goal cell is active at the goal site (rectangular box). Synaptic modification between a subicular and a goal cell is only allowed if the goal cell is active and the subicular cell fires in the late phase of theta (shaded area). Here, subicular cells 4 and 5 fire late in the cycle of theta when the rat is near the goal. **c** Firing rate of an 'East' goal cell after learning. As the only synapses strengthened were from subicular cells 4, 5, the 'East' goal cell will only be active when the rat is within the unified place fields of subicular cells 4 and 5 which are located predominantly east to the goal (see also subfigure A, shaded area). At the overlap of these two place fields the goal cell receives higher excitation (input from both cells) causing the firing rate tuning curve of the 'East' goal cell to be a unimodal curve peaking East of the goal (dashed line shows location of the goal area)

distinguishing between predictive and inertial firing, as the plasticity between goal and place cells did not take firing phase into account. Instead, formation of proper connections from place cells to goal cells relied on an appropriate connectivity matrix within CA3, which in turn depended on reliable plasticity between CA3 cells. Thus, the importance of phase precession lies in ensuring that CA3 neurons having place fields along the trajectory of the animal emitted spikes repeatedly and consistently in the appropriate order and in a compressed form with appropriate timing for LTP to occur (Fig. 1B). This role is essentially the same as that played in one-shot episodic learning in the model of Yamaguchi (2003) and Sato and Yamaguchi (2003), discussed above.

4 Discussion

Based on the theories described above, theta oscillations appear to be versatile, able to aid neural computations in many different ways. Most of these models are appealing as they offer a direct link between theta and improved performance in either some fundamental neural computation or in some specific behavioral task. However, despite the wealth of models, there are also some troubling issues that make it hard to synthesize the knowledge we gained through theoretical work about the function of theta oscillations.

There are over a dozen models proposing different functions for theta. On one hand, this may be greeted with enthusiasm as an indication of the fundamental importance of theta oscillations in neural computations. On the other hand, this diversity may be puzzling because different models are often incompatible with each other. For example, models that assume limit cycle oscillations in the firing rates of neurons (e.g., Scarpetta et al. 2002) are hard to reconcile with those assuming exactly one spike per cell per theta cycle (e.g. Hopfield 1995; Brody and Hopfield 2003) especially when they all claim to account for the same phenomenon (odor recognition in the olfactory bulb in this example). Which model should we believe?

The 'objective' way to choose between different models is to test their predictions. Unfortunately, many of the models discussed here do not make direct, easily testable predictions. In particular, one type of prediction that is often made in models on the function of theta compares the performance of the network with and without theta and demonstrates an impairment in the latter case. This type of prediction is genuinely hard to test, because the theta rhythm itself is already an emergent property of neural network dynamics. Thus, experimental manipulations that affect theta power are likely to influence network dynamics in many ways and it is hard to assess whether their effects on cognitive performance is causally related to their effect on theta, or it is just a correlation due to common factors.

In one study, for example, application of a local anesthetic in the septum of freely behaving rats abolished hippocampal theta oscillations while still preserving place-specific cell firing in CA1, the main output area of the hippocampus, and resulted in impairment in a spatial working memory task (Mizumori et al. 1989). This experiment seems to control for indirect effects of the manipulation (by demonstrating preserved place cell firing) and may provide direct support for a specific role of theta in hippocampal function, and therefore for models that assign theta such a function. Yet, even in this study, it is controversial whether the behavioral impairment is only due to the lack of hippocampal theta or the septally applied anesthetic also had other side effects that could not be controlled for by the experimenters (e.g., a substantial change in hippocampal interneuronal firing rates and in CA3 place cell firing was also found). Furthermore, in this study only a gross measure of behavioral performance (error rates) was published, yet different models could, or at least should

make different specific predictions about the exact nature of the impairment that is caused by 'switching off' theta.

There is also a more conceptual difficulty with comparing the 'theta on' and 'theta off' versions of the same model. This comparison is often made to demonstrate from a theoretical point of view why theta is necessary for some specific neural computation, i.e., why it is better to have theta than not to have it. However, as we have already pointed out in the case of some models, to answer such a question by making the fair comparison would require implementing the same computation in two independent models, one but not the other using theta oscillations, and show that some fundamental limitations (such as demonstrated by Li and Dayan (1999)) can only be transcended with the use of oscillatory dynamics. These comparisons are admittedly hard to make, because the same function can be implemented in any number of different ways, and it can hardly be guaranteed that from all possible 'theta off' models the best one was chosen. Still, in cases when there exists a well-defined optimality criterion this enterprise may be possible and therefore fruitful.

After all, it may also very well be that a unified model of theta function is not needed, because theta oscillations in different areas merely coincidentally fall into the same frequency band, but are fundamentally different in terms of the neural dynamics they entail and the neural computations they support. Even theta in the same brain area but under different behavioral modes, such as hippocampal theta during REM sleep and awake behavior, may serve different functions. Although such coincidences may not be ruled out, they should be treated with scepticism, especially in the light of some classical (Macrides et al. 1982) and a number of more recent findings (Sarnthein et al. 1998; Stam et al. 2002; Fell et al. 2003; Mizuhara et al. 2004; Seidenbecher et al. 2003; Sauseng et al. 2004) from various species and brain regions about increased coherence of theta oscillations between different brain areas during some cognitively demanding tasks. These results seem to support a more unified function of theta oscillations. Notably, with a few exceptions (e.g., Jensen 2001), there are no computational theories to date about how such an oscillatory interaction of different brain areas could aid neural computations.

Finally, we would like to highlight another area of research where we expect theta-related models to make important contributions. In the past years, increasing evidence has accumulated as to the significance of spontaneous ongoing network activity in different brain regions (Harris et al. 2003; Kenet et al. 2003; Fiser et al. 2004). These experimental studies show that changes in neuronal activities cannot be accounted for by external variables, such as visual input or the position of the animal, that were traditionally thought to control the firing of orientation-selective V1 neurons (Kenet et al. 2003; Fiser et al. 2004) or hippocampal place cells (Harris et al. 2003). Instead, external variables have now been shown to have subtler, modulatory effects on spontaneously emerging patterns of neural network activity. Quite remarkably, these spontaneous patterns of activity can be recorded even in the total absence of environmental inputs and often

take the form of oscillations, of which the frequency at least in some cases fall into the theta band (Fiser et al. 2004). As most current theories and models simply predict zero or unpatterned activity in the absence of external inputs, these results seem to fundamentally challenge the often adopted but clearly unsatisfactory view of neural networks as mere input-output devices. Analyzing the functional role of theta oscillations in these cases may thus allow us to better understand the neurobiological processes underlying perception and cognition.

Acknowledgements The authors are especially grateful to Peter Dayan for many useful discussions and comments on the manuscript. We also thank Ole Jensen for many useful comments on the manuscript and Gergő Orbán for discussions on an earlier version of this work. M.L. was supported by the Gatsby Charitable Foundation, Zs.H. and P.É. were supported by Hungarian Scientific Research Fund (OTKA) grant T038140 and by the Henry R. Luce Foundation.

References

- Adrian ED (1942) Olfactory reactions in the brain of the hedgehog *J Physiol* 100:459–473
- Alonso A, Garcia-Austt E (1987) Neuronal sources of theta rhythm in the entorhinal cortex of the rat. II. Phase relations between unit discharges and theta field potentials. *Exp Brain Res* 67:502–509
- Alonso A, Klink R (1993) Differential electroresponsiveness of stellate and pyramidal-like cells of medial entorhinal cortex layer II. *J Neurophysiol* 70:128–143
- Amit DJ (1989) Modeling brain function. Cambridge University Press, Cambridge
- Arbib MA, Erdi P, Szentágothai J (1998) Neural organization: structure, function, and dynamics. MIT Press, Cambridge
- August DA, Levy WB (1999) Temporal sequence compression by an integrate-and-fire model of hippocampal area CA3. *J Comput Neurosci* 6:71–90
- Baird B (1986) Nonlinear dynamics of pattern formation and pattern recognition in the rabbit olfactory bulb. *Physica D* 22:150–175
- Baird B (1990) Bifurcation and learning in network models of oscillating cortex. *Physica D* 42:365–384
- Bi GQ, Poo MM (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and post-synaptic cell type. *J Neurosci* 18:10464–10472
- Bi G, Poo MM (2001) Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu Rev Neurosci* 24:139–166
- Bogacz R, Brown MW, Giraud-Carrier C (2001) Model of familiarity discrimination in the perirhinal cortex. *J Comput Neurosci* 10:5–23
- Borisjuk RM, Hoppensteadt FC (1998) Memorizing and recalling spatial-temporal patterns in an oscillator model of the hippocampus. *Biosystems* 48:3–10
- Borisjuk RM, Hoppensteadt FC (1999) Oscillatory models of the hippocampus: a study of spatio-temporal patterns of neural activity. *Biol Cybern* 81:359–371
- Borisjuk RM, Denham M, Hoppensteadt FC, Kazanovich Y, Vinogradova O (2000) An oscillatory neural network model of sparse distributed memory and novelty detection. *Biosystems* 58:265–272
- Brankack J, Stewart M, Fox S E (1993) Current source density analysis of the hippocampal theta rhythm: associated sustained potentials and candidate synaptic generators. *Brain Res* 615:310–327
- Brody CD, Hopfield JJ (2003) Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron* 37:843–852
- Burgess N, Recce M, O'Keefe J (1994) A model of hippocampal function. *Neural Netw* 7:1065–1081

- Csicsvári J, Hirase H, Czurkó A, Mamiya A, Buzsáki Gy (1999) Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *J Neurosci* 19:274–287
- Dragoi G, Harris KD, Buzsáki Gy (2003) Place representation within hippocampal networks is modified by long-term potentiation. *Neuron* 39:843–853
- Duda RO, Hart PE, Stork DG (2001) *Pattern classification*. 2nd ed, Wiley, New York
- Engel AK, Fries P, Singer W (2003) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 4:21:844–848
- Ermentrout GB, Cowan JD (1979) A mathematical theory of visual hallucination patterns. *Biol Cybern* 34:137–150
- Fell J, Klaver P, Elfadil H, Schaller C, Elger C E, Fernandez G (2003) Rhinal-hippocampal theta coherence during declarative memory formation: interaction with gamma synchronization? *Eur J Neurosci* 17:1082–1088
- Fiser J, Chiu C, Weliky M (2004) Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431:573–578
- Frank L M, Brown E N, Wilson M A (2001) A comparison of the firing properties of putative excitatory and inhibitory neurons from CA1 and the entorhinal cortex. *J Neurophysiol* 86:2029–2040
- Freund T F, Buzsáki Gy (1996) Interneurons of the hippocampus. *Hippocampus* 6:347–470
- Gevins A, Smith ME, McEvoy L, Yu D (1997) High-resolution EEG mapping of cortical activation related to working memory: effects of task difficulty, type of processing, and practice. *Cereb Cortex* 7:374–385
- Gray CM (1994) Synchronous oscillations in neuronal systems: mechanisms and functions. *J Comput Neurosci* 1:11–38
- Green J, Arduini A (1954) Hippocampal electrical activity in arousal. *J Neurophysiol* 17:533–557
- Hajós M, Hoffmann WE, Orbán G, Kiss T, Érdi P (2004) Modulation of septo-hippocampal θ activity by GABA_A receptors: an experimental and computational approach. *Neuroscience* 126:599–610
- Harris KD, Csicsvári J, Hirase H, Dragoi G, Buzsáki Gy (2003) Organization of cell assemblies in the hippocampus. *Nature* 424:552–556
- Hasselmo ME, Schnell E, Barkai E (1995) Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region CA3. *J Neurosci* 15:5249–5262
- Hasselmo ME, Wyble BP, Wallenstein GV (1996) Encoding and retrieval of episodic memories: role of cholinergic and GABAergic modulation in the hippocampus. *Hippocampus* 6:693–708
- Hasselmo ME, Bodelon C, Wyble BP (2002) A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput* 14:793–817
- Hasselmo ME, Hay J, Ilyin M, Gorchetnikov A (2002) Neuromodulation, theta rhythm and rat spatial navigation. *Neural Netw* 15:689–707
- Hinton GE (2002) Training products of experts by minimizing contrastive divergence. *Neural Comput* 14:1171–1800
- Hölscher C, Anwyl R, Rowan MJ (1997) Stimulation on the positive phase of hippocampal theta rhythm induces long-term potentiation that can be depotentiated by stimulation on the negative phase in area CA1 *in vivo*. *J Neurosci* 17:6470–6477
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 76:2554–2558
- Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–36
- Hopfield JJ (2004) Encoding for computation: recognizing brief dynamical patterns by exploiting effects of weak rhythms on action-potential timing. *Proc Natl Acad Sci USA* 101:6255–6260
- Huerta PT, Lisman JE (1995) Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 *in vitro*. *Neuron* 15:1053–1063
- Jensen O (2001) Information transfer between rhythmically coupled networks: reading the hippocampal phase code. *Neural Comput* 13:2743–2761
- Jensen O, Lisman JE (1996a) Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learn Mem* 3:279–287
- Jensen O, Lisman JE (1996b) Novel lists of 7 ± 2 known items can be reliably stored in an oscillatory short-term memory network: interaction with long-term memory. *Learn Mem* 3:257–263
- Jensen O, Lisman JE (1996c) Theta/gamma networks with slow NMDA channels learn sequences and encode episodic memory: role of NMDA channels in recall. *Learn Mem* 3:264–278
- Jensen O, Lisman JE (1998) An oscillatory short-term memory buffer model can account for data on the Sternberg task. *J Neurosci* 18:10688–10699
- Jensen O, Lisman J (2005) Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends Neurosci* 28:67–72
- Jensen O, Tesche CD (2002) Frontal theta activity in humans increases with memory load in a working memory task. *Eur J Neurosci* 15:1395–1399
- Jensen O, Idiart M, Lisman JE (1996) Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: role of fast NMDA channels. *Learn Mem* 3:243–256
- Jouvet M, Michel F, Courjon J (1959) Sur un stade d'activité électrique cérébrale rapide au cours du sommeil physiologique. *C R Seances Soc Biol Fil* 153:1024–1028
- Jung R, Kornmüller A (1938) Ein methodik der abteilung lokalisierter potential schwankungen aus subcorticalen hirnyebieten. *Arch Psychiat Neuroenkr* 109:1–30
- Kahana MJ, Sekuler R, Caplan JB, Kirschen M, Madsen JR (1999) Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature* 399:781–784
- Kahana MJ, Seelig D, Madsen JR (2001) Theta returns. *Curr Opin Neurobiol* 11:739–744
- Káli Sz, Dayan P (2004) Off-line replay maintains declarative memories in a model of hippocampal-neocortical interactions. *Nat Neurosci* 7:286–294
- Kamondi A, Acsády L, Wang XJ, Buzsáki Gy (1998) Theta oscillation in somata and dendrites of hippocampal pyramidal cells *in vivo*: activity-dependent phase-precession of action potentials. *Hippocampus* 8:244–261
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425:954–956
- King C, Recce M, O'Keefe J (1998) The rhythmicity of cells of the medial septum/diagonal band of Broca in the awake freely moving rat: relationships with behaviour and hippocampal theta. *Eur J Neurosci* 10:464–477
- Klausberger T, Magill PJ, Márton LF, Roberts JD, Cobden PM, Buzsáki Gy, Somogyi P (2003) Brain-state- and cell-type-specific firing of hippocampal interneurons *in vivo*. *Nature* 421:844–848
- Klimesch W (1999) EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Res Rev* 29:169–195
- Klink R, Alonso A (1993) Ionic mechanisms for the subthreshold oscillations and differential electroresponsiveness of medial entorhinal cortex layer II neurons. *J Neurophysiol* 70:144–157
- Koene RA, Gorchetnikov A, Cannon RC, Hasselmo ME (2003) Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. *Neural Netw* 16:577–584
- Latham PE, Deneve S, Pouget A (2003) Optimal computation with attractor networks. *J Physiol Paris* 97:683–694
- Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36:1183–1194
- Lee H, Simpson GV, Logothetis NK, Rainer G (2005) Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron* 45:147–156
- Legenstein R, Markram H, Maas W (2003) Input prediction and autonomous movement analysis in recurrent circuits of spiking neurons. *Rev Neurosci* 14:5–19

- Lengyel M, Dayan P (2005) Rate- and phase-coded autoassociative memory. *Advances in neural information processing systems* 17 (in press)
- Lengyel M, Érdi P (2004) Theta modulated feed-forward network generates rate and phase coded firing in the entorhino-hippocampal system. *IEEE Trans Neural Netw* 15:1092–1099
- Leung LS, Yim CY (1986) Intracellular records of theta rhythm in hippocampal CA1 cells of the rat. *Brain Res* 367:323–327
- Leung LS, Yu HW (1998) Theta-frequency resonance in hippocampal CA1 neurons in vitro demonstrated by sinusoidal current injection. *J Neurophysiol* 79:1592–1596
- Levy WB (1996) A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* 6:579–590
- Li Z, Dayan P (1999) Computational differences between asymmetrical and symmetrical networks. *Network* 10:59–77
- Li Z, Hopfield JJ (1989) Modeling the olfactory bulb and its neural oscillatory processings. *Biol Cybern* 61:379–392
- Lisman JE, Idiart MAP (1995) Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* 267:1512–1515
- Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron* 29:145–156
- Maass W, Natschläger T, Markram H (2002) Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput* 14:2531–2560
- Macrides F, Eichenbaum HB, Forbes WB (1982) Temporal relationship between sniffing and the limbic theta rhythm during odor discrimination reversal learning. *J Neurosci* 2:1705–1717
- Maei HR, Latham PE (2004) How can realistic networks process time-varying signals? *Soc Neurosci (Abst)* 30: 81.9
- Margrie TW, Schaefer AT (2003) Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J Physiol* 546:363–74
- Markram H, Lubke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic apss and epsps. *Science* 275:213–215
- Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci USA* 94:8918–8921
- Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417:741–746
- Mitman KE, Laurent PA, Levy WB (2003) Defining time in a minimal hippocampal CA3 model by matching time-span of associative synaptic modification and input pattern duration. In: *Proceedings of the international joint conference on neural networks* pp 1631–1636
- Mizuhara H, Wang LQ, Kobayashi K, Yamaguchi Y (2004) A long-range cortical network emerging with theta oscillation in a mental task. *Neuroreport* 15:1233–1238
- Mizumori SJY, McNaughton BL, Barnes CA, Fox KB (1989) Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of CA3c output: evidence for pattern completion in hippocampus. *J Neurosci* 9:3915–3928
- Myers CE, Ermita BR, Harris K, Hasselmo M, Solomon P, Gluck MA (1996) A computational model of cholinergic disruption of septo-hippocampal activity in classical eyeblink conditioning. *Neurobiol Learn Mem* 66:51–66
- Nádasy Z, Hirase H, Czurkó A, Csicsvári J, Buzsáki Gy (1999) Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci* 19:9497–9507
- Nishikawa T, Lai YC, Hoppensteadt FC (2004) Capacity of oscillatory associative-memory networks with error-free retrieval. *Phys Rev Lett* 92:108101
- Núñez A, Garcia-Austt E, Buño Jr W (1987) Intracellular theta-rhythm generation in identified hippocampal pyramids. *Brain Res* 416:289–300
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Res* 34:171–175
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Clarendon Press, Oxford
- O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330
- Orr G, Rao G, Houston FP, McNaughton BL, Barnes CA (2001) Hippocampal synaptic plasticity is modulated by theta rhythm in the fascia dentata of adult and aged freely behaving rats. *Hippocampus* 11:647–654
- Paulsen O, Sejnowski TJ (2000) Natural patterns of activity and long-term synaptic plasticity. *Curr Opin Neurobiol* 10:172–179
- Pavrides C, Greenstein YJ, Grudman M, Winson J (1988) Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of theta-rhythm. *Brain Res* 439:383–387
- Pike FG, Goddard RS, Suckling JM, Ganter P, Kasthuri N, Paulsen O (2000) Distinct frequency preferences of different types of rat hippocampal neurones in response to oscillatory input currents. *J Physiol* 529 Pt 1:205–213
- Pouget A, Dayan P, Zemel RS (2003) Inference and computation with population codes. *Annu Rev Neurosci* 26:381–410
- Raghavachari S, Kahana MJ, Rizzuto DS, Caplan JB, Kirschen MP, Bourgeois B, Madsen JR, Lisman JE (2001) Gating of human theta oscillations by a working memory task. *J Neurosci* 21:3175–3183
- Rinzel J, Ermentrout B (1998) Analysis of neural excitability and oscillations. In: Koch C, Segev I (eds) *Methods in neuronal modeling*, A Bradford Book. MIT Press, Cambridge, pp 251–291
- Rosenblatt F (1958) The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev* 65:386–408
- Sarnthein J, Petsche H, Rappelsberger P, Shaw GL, von Stein A (1998) Synchronization between prefrontal and posterior association cortex during human working memory. *Proc Natl Acad Sci USA* 95:7092–7096
- Sato N, Yamaguchi Y (2003) Memory encoding by theta phase precession in the hippocampal network. *Neural Comput* 15:2379–2397
- Sauseng P, Klimesch W, Doppelmayr M, Hanslmayr S, Schabus M, Gruber WR (2004) Theta coupling in the human electroencephalogram during a working memory task. *Neurosci Lett* 354:123–126
- Scarpetta S, Li Z, Hertz J (2002) Hebbian imprinting and retrieval in oscillatory neural networks. *Neural Comput* 14:2371–2396
- Seidenbecher T, Laxmi TR, Stork O, Pape HC (2003) Amygdalar and hippocampal theta rhythm synchronization during fear memory retrieval. *Science* 301:846–850
- Skaggs WE, McNaughton BL (1996) Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271:1870–1873
- Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6:149–172
- Skarda CA, Freeman WJ (1987) How brains make chaos in order to make sense of the world. *Behav Brain Sci* 10:161–173
- Sohal VS, Hasselmo M (1998a) GABA_B modulation improves sequence disambiguation in computational models of hippocampal region CA3. *Hippocampus* 8:171–193
- Sohal VS, Hasselmo ME (1998b) Changes in GABA_B modulation during a theta cycle may be analogous to the fall of temperature during simulated annealing. *Neural Comput* 10:869–882
- Stam CJ, van Cappellen van Walsum S, Micheloyannis S (2002) Variability of EEG synchronization during a working memory task in healthy subjects. *Int J Psychophysiol* 46:53–66
- Tesche CD, Karhu J (2000) Theta oscillations index human hippocampal activation during a working memory task. *Proc Natl Acad Sci USA* 97:919–924
- Trullier O, Meyer JA (2000) Animat navigation using a cognitive graph. *Biol Cybern* 83:271–285
- Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL (1996) Population dynamics and theta phase precession of hippocampal place cell firing: a spiking neuron model. *Hippocampus* 6:271–280
- Tsuda I (2001) Toward an interpretation of dynamic neural activity in terms of chaotic dynamical systems. *Behav Brain Sci* 24:703–810

- Uchida S, Nishida M, Hirai N, Miwakeichi F (2001) Does hippocampal theta exist in the human brain? *Neurobiol Sleep-Wakefulness Cycl* 1:1–8
- Vanderwolf CH (1969) Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol* 26:407–418
- Wallenstein GV, Hasselmo ME (1997) GABAergic modulation of hippocampal activity: sequence learning, place field development, and the phase precession effect. *J Neurophysiol* 78:393–408
- Wallenstein GV, Eichenbaum H, Hasselmo ME (1998) The hippocampus as an associator of discontinuous events. *Trends Neurosci* 21:317–323
- Wang DL, Buhmann J, von der Malsburg C (1990) Pattern segmentation in associative memory. *Neural Comput* 2:94–384
- Wills SA (2004) Computation with spiking neurons. PhD Thesis, University of Cambridge
- Wilson HR, Cowan JD (1972) Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys J* 12:1–24
- Wyble BP, Linster C, Hasselmo ME (2000) Size of CA1-evoked synaptic potentials is related to theta rhythm phase in rat hippocampus. *J Neurophysiol* 83:2138–2144
- Yamaguchi Y (2003) A theory of hippocampal memory based on theta phase precession. *Biol Cybern* 89:1–9
- Ylinen A, Soltész I, Bragin A, Penttonen M, Sík A, Buzsáki Gy (1995) Intracellular correlates of hippocampal theta rhythm in identified pyramidal cells, granule cells, and basket cells. *Hippocampus* 5:78–90