Synaptic learning rules for sequence learning

Eric T. Reifenstein[1,2,3] and Richard Kempter[1,2,3]

[1]Institute for Theoretical Biology, Department of Biology, Humboldt-Universität zu Berlin, Philippstr. 13, D-10115 Berlin

[2]Bernstein Center for Computational Neuroscience Berlin, Philippstr. 13, D-10115 Berlin

[3]Einstein Center for Neurosciences Berlin, Charitéplatz 1, D-10117 Berlin

Abstract

Remembering the temporal order of a sequence of events is a task easily performed by humans in everyday life, but the underlying neuronal mechanisms are unclear. This problem is particularly intriguing as human behavior often proceeds on a time scale of seconds, which is in stark contrast to the much faster millisecond time-scale of neuronal processing in our brains. One long-held hypothesis in sequence learning suggests that a particular temporal fine-structure of neuronal activity — termed "phase precession" — enables the compression of slow behavioral sequences down to the fast time scale of the induction of synaptic plasticity. Using mathematical analysis and computer simulations, we find that — for short enough synaptic learning windows — phase precession can improve temporal-order learning tremendously and that the asymmetric part of the synaptic learning window is essential for temporal-order learning. To test these predictions, we suggest experiments that selectively alter phase precession or the learning window and evaluate memory of temporal order.

Introduction

It is a pivotal quality for animals to be able to store and recall the order of events ("temporal-order learning", 1–3) but there is only little work on the neural mechanisms generating asymmetric memory associations across behavioral time intervals (4). Putative mechanisms need to bridge the gap between the faster time scale of the induction of synaptic plasticity (typically milliseconds) and the slower time scale of behavioral events (seconds or slower). The slower time scale of behavioral events is mirrored, for example, in the time course of firing rates of hippocampal place cells (5), which signal when an animal visits certain locations ("place fields") in the environment. The faster time scale is given by the temporal properties of the induction of synaptic plasticity (6, 7) — and spike-timing dependent plasticity (STDP) is a common form of synaptic plasticity that depends on the millisecond timing and temporal order of presynaptic and postsynaptic spiking. For STDP, the so-called "learning window" describes the temporal intervals at which presynaptic and postsynaptic activity induce synaptic plasticity. Such precisely timed neural activity can be generated by phase precession, which is the successive across-cycle shift of spike phases from late to early with respect to a background oscillation (Fig. 1). As an animal explores an environment, phase precession can be observed in the activity of hippocampal place cells with respect to the theta oscillation (8, 9). Phase precession is highly significant in single trials (10, 11) and occurs even in first traversals of a place field in a novel environment (12). Interestingly, phase precession allows for a temporal compression of a sequence of behavioral events from the time scale of seconds down to milliseconds (Fig. 1; 12–14), which matches the widths of generic STDP learning windows (15–18). This putative advantage of phase precession for temporal-order learning, however, has not yet been quantified. To assess the benefit of phase precession for temporal-order learning, we determine the synaptic weight change between pairs of cells whose activity represents two events of a sequence. Using both analytical methods and numerical simulations, we find that phase precession can dramatically facilitate temporal-order learning by increasing the synaptic weight change and the signal-to-noise ratio by up to an order of magnitude. We thus provide a mechanistic description of associative chaining models (19) and extend these models to explain how to store serial order.

Results

To address the question of how behavioral sequences could be encoded in the brain, we study the change of synapses between neurons that represent events in a sequence. We assume that the temporal order of two events is encoded in the asymmetry of the efficacies of synapses that connect neurons representing the two events (Fig. 1). After the successful encoding of a sequence, a neuron that was activated earlier in the sequence has a strengthened connection to a neuron that was activated later in the sequence; whereas the connection in the reverse direction may be unchanged or is even weakened. As a result, when the first event is encountered and/or the first neuron is activated, the neuron representing the second event is activated. Consequently, the behavioral sequence could be replayed (as illustrated by simulations e.g., in 14, 20–26) and the memory of the temporal order of events is recalled (27, 28). We note, however, that in what follows we do not simulate such a replay of sequences, which would depend also on a vast number of parameters that define the network; instead, we rather focus on the underlying change in connectivity, which is the very basis of replay, and draw connections to "replay" in the *Discussion*.

Let us now illustrate key features of the encoding of the temporal order of sequences. To do so, we consider the weight change induced by the activity of two sequentially activated cells i and j that represent two behavioral events (dashed lines in Fig. 2A). Classical Hebbian learning (29), where weight changes Δw_{ij} depend on the product of the firing rates f_i and f_j, is not suited for temporal-order learning because the weight change is independent of the order of cells:

$$\Delta w_{ij} \propto f_i \cdot f_j = f_j \cdot f_i \propto \Delta w_{ji} \,.$$

Therefore, a classical Hebbian weight change is symmetric, i.e., $\Delta w_{ij} - \Delta w_{ji} = 0$. This result can be generalized to learning rules that are based on the product of two arbitrary functions of the firing rates. We note that, although not suited for temporal-order learning, Hebbian rules are able to achieve more general "sequence learning", where an association between sequence elements is created — independent of the order of events. To become sensitive to temporal order, we use spike-timing dependent plasticity (STDP; 6, 7). For STDP, average

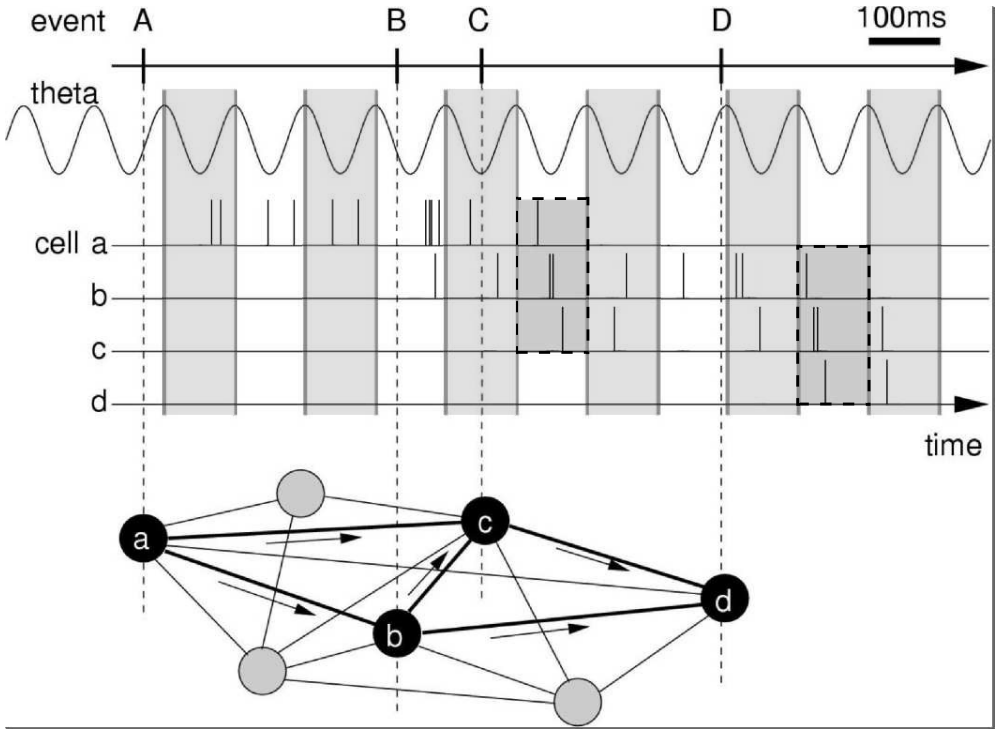

Figure 1: Rationale for temporal-order learning via phase precession. Top: Behavioral events (A to D) happen on a time scale of seconds. Middle: These events are represented by different cells (a-d), which fire a burst of stochastic action potentials in response to the onset of their respective event. We assume that each cell shows phase precession with respect to the LFP's theta oscillation (every second cycle is marked by a grey box). When the activities of multiple cells overlap, the sequence of behavioral events is compressed in time to within one theta cycle (two examples highlighted in the dashed, shaded boxes). **Bottom**: This faster time scale can be picked up by STDP and strengthen the connections between the cells of the sequence.

weight changes depend on the cross-correlation function of the firing rates (example in Fig. 2C,D),

$$C_{ij}(t) := \int_{-\infty}^{\infty} \mathrm{d}t' \, f_i(t') f_j(t'+t) \,,$$

which is anti-symmetric: $C_{ij}(t) = C_{ji}(-t)$. Assuming additive STDP, i.e., weight changes resulting from pairs of pre- and postsynaptic action potentials are added, the average synaptic weight change Δw_{ij} between the two cells in a sequence can then be calculated explicitly (30):

$$\Delta w_{ij} = \int_{-\infty}^{+\infty} \mathrm{d}t \, W(t) C_{ij}(t) \tag{1}$$

where W is the STDP learning window (example in Fig. 2E). We aim solve equation 1 for given firing rates f_i

and f_j. To do so, we assume that the synaptic weight w_{ij} is generally small and thus only has a weak impact on the cross-correlation of the cells during encoding, that is, for the "encoding" of a sequence the cross-correlation function is dominated by feedforward input whereas the recurrent inputs are neglected.

Next, let us show that the symmetry of W is essential for temporal-order learning. Any learning window W can be split up into an even part W^{even}, with $W^{\text{even}}(t) = W^{\text{even}}(-t)$, and an odd part W^{odd}, with $W^{\text{odd}}(t) = -W^{\text{odd}}(-t)$, such that $W = W^{\text{even}} + W^{\text{odd}}$. For even learning windows, one can derive from Equation 1 and the anti-symmetry of C_{ij} that weight changes are symmetric, i.e. $\Delta w_{ij} = \Delta w_{ji}$; therefore only the odd part W^{odd} of W is useful for learning temporal order.

To further explore requirements for encoding the temporal order of a sequence of events, we restrict our analysis to odd learning windows. We then can relate the weight change Δw_{ij} to the essential features of $C_{ij}(t)$. To do so, we integrate Equation 1 by parts (with W replaced by W^{odd}),

$$\Delta w_{ij} = \underbrace{\left[\overline{W^{\text{odd}}}(t) \cdot C_{ij}(t)\right]_{-\infty}^{+\infty}}_{=0} + \int_{-\infty}^{+\infty} \mathrm{d}t \, \left[-\overline{W^{\text{odd}}}(t)\right] \, C'_{ij}(t) \,, \tag{2}$$

with the primitive $\overline{W^{\text{odd}}}(t) := \int_{-\infty}^{t} \mathrm{d}t' \, W^{\text{odd}}(t')$ and the derivative $C'_{ij}(t) := \frac{\mathrm{d}}{\mathrm{d}t} C_{ij}(t)$. Because $\overline{W^{\text{odd}}}(t)$ can be assumed to have finite support (note that $\int_{-\infty}^{+\infty} \mathrm{d}t \, W^{\text{odd}}(t) = 0$), the first term in Equation 2 vanishes. Also the learning window has finite support, and therefore we can restrict the integral in the second term in Equation 2 to a finite region of width K around zero:

$$\Delta w_{ij} = \int_{|t|<K} \mathrm{d}t \, \left[-\overline{W^{\text{odd}}}(t)\right] \, C'_{ij}(t) \tag{3}$$

where K describes the width of the learning window W (gray region in Fig. 2E). The integral in Equation 3 can be interpreted as the cross-correlation's slope around zero, weighted by the symmetric function $-\overline{W^{\text{odd}}}(t)$; interestingly, features of C_{ij} for $|t| \gg K$, for example whether side lobes of the correlation function are decreasing or not, are irrelevant.

As a generic example of sequence learning, let us consider the activities of two cells i and j that encode two behavioral events, for example the traversal of two place fields of two hippocampal place cells. In general, the cells' responses to these events are called "firing fields". We model these firing fields as two Gaussian functions $G_{0,\sigma}$ and $G_{T_{ij},\sigma}$ that have the same width σ but different mean values 0 and T_{ij} (we note that T_{ij} and σ are

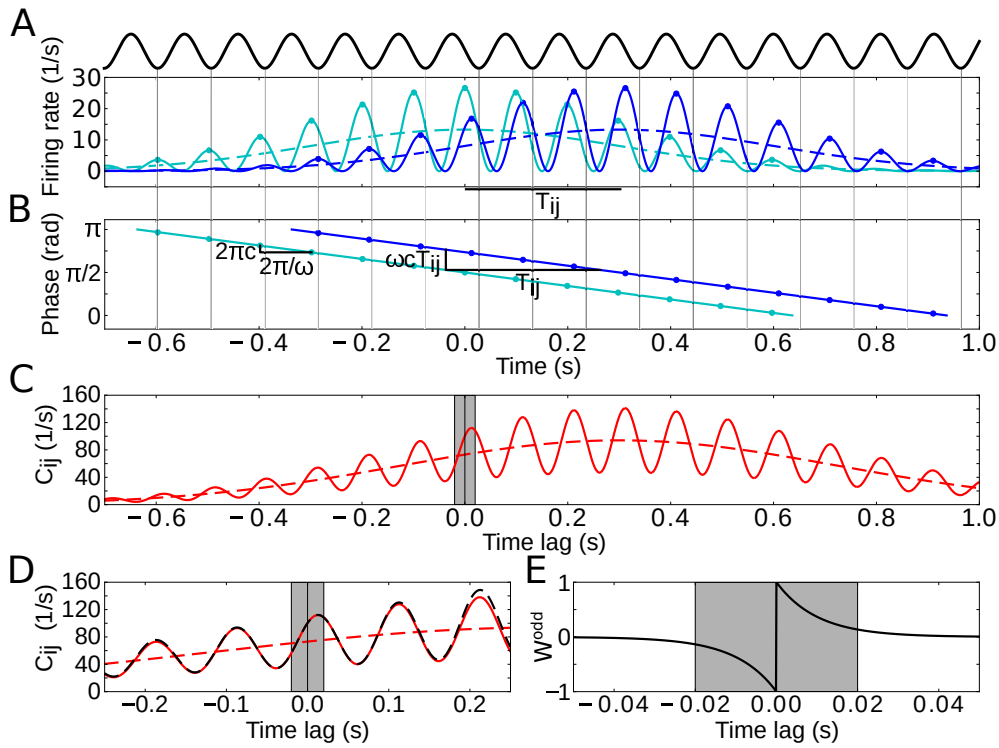

Figure 2: Model of two sequentially activated phase-precessing cells. **(A)** Oscillatory firing-rate profiles for two cells (solid blue and cyan lines). The black curve depicts the population theta oscillation. For easier comparison of the two different frequencies, the population activity's troughs are continued by thin gray lines, and the peaks of the cell-intrinsic theta oscillation are marked by dots. Dashed lines depict the underlying Gaussian firing fields without theta modulation. **(B)** Phase precession of the two cells (same colors as in A). The compression factor c describes the phase shift per theta cycle for an individual cell ($2\pi c$). For the temporal separation T_{ij} of the firing fields and the theta frequency ω, the phase difference between the cells is $\omega c T_{ij}$. The dots depict the times of the maxima in (A). **(C)** Resulting cross-correlation for the two firing rates from (A). The solid red curve shows the full cross-correlation. The dashed line depicts the cross-correlation without theta-modulation. The grey region indicates small (< 20 ms) time lags. **(D)** Same as in (C), but zoomed in. Note that the first peak of the theta modulation is at a positive non-zero time lag, reflecting phase precession. The dashed black curve shows the approximation of the cross-correlation for the analytical treatment (Materials and Methods, Equation 17). **(E)** Synaptic learning window. The grey region indicates the region in which the learning window is large, and this region is also indicated in (C) and (D). Positive time lags correspond to postsynaptic activity following presynaptic activity. Parameters for all plots: $T_{ij} = 0.3$ s, $\omega = 2\ \pi \cdot 10$ Hz, $\sigma = 0.3$ s, $\tau = 10$ ms, $\mu = 1$, $c = 0.042$, $A = 10$.

measured in units of time, i.e., seconds; Fig. 2A, dashed curves). In this case of identical Gaussian shapes of the two firing fields, the cross-correlation $C_{ij}(t)$ is also a Gaussian function, denoted by $G_{T_{ij},\sqrt{2}\sigma}$, but with mean T_{ij} and width $\sqrt{2}\sigma$ (dashed curve in Fig. 2C). The value $\sigma = 0.3$ s, which we use in the example of Fig. 2, matches experimental findings on place cells (8, 31).

It is widely assumed that phase precession facilitates temporal-order learning (10, 13, 32), but it has never been quantitatively shown. To test this hypothesis and to calculate how much phase precession contributes to temporal-order learning, we consider Gaussian firing fields that exhibit oscillatory modulations with theta frequency ω (Fig. 2A, solid curves). The time-dependent firing rate of cell i is described by $f_i(t) \propto G_{\mu_i,\sigma}(t)\,\{1 + \cos[\omega(t - c\mu_i)]\}$, i.e., a Gaussian that is multiplied by a sinusoidal oscillation; see also Equation 11 in *Materials and Methods*. Phase precession occurs with respect to the population theta, which oscillates at a frequency of $(1-c)\omega$ that is slightly smaller than ω, with a "compression factor" c that is usually small: $0 \leq c \ll 1$ (31, 32). This compression factor c describes the average advance of the firing phase — from theta cycle to theta cycle — in units of the fraction of a theta cycle; c thus determines the slope ωc of phase precession (Fig. 2B). A typical value is $c \approx \pi/(4\sigma\omega)$, which accounts for "slope-size matching" of phase precession (31); that is, c is inversely proportional to the field size $L := 4\sigma$ of the firing field, and the total range of phase precession within the firing field is constant and equals $\pi \equiv 180°$. If there are multiple theta oscillation cycles within a firing field ($\omega\sigma \gg 1$), which is typical for place cells, the cross-correlation $C_{ij}(t)$ is a theta modulated Gaussian (solid curve in Fig. 2C; see also Equation 15 in *Materials and Methods*).

The generic shape of the cross-correlation C_{ij} in Fig. 2C allows for an advanced interpretation of Equation 3, which critically depends on the width K of the learning window W. We distinguish here two limiting cases: narrow learning windows ($K \ll 1/\omega \ll \sigma$), i.e., the with K of the learning window is much smaller than a theta cycle and the width of a firing field, and wide learning windows ($K \gg \sigma$), i.e., the width K of the learning window exceeds the width of a firing field. Let us first consider narrow learning windows. Only later in this manuscript, we will turn to the case of wide learning windows.

Dependence of temporal-order learning on the overlap of firing fields for narrow learning windows ($K \ll 1/\omega \ll \sigma$)

We first show formally that sequence learning with narrow learning windows requires that the two firing fields do overlap, i.e., their separation T_{ij} should be less than or at least similar to the width σ of the firing fields. In Equation 3, which was derived for odd learning windows, the weight change Δw_{ij} is determined by $C'_{ij}(t)$ around $t = 0$ in a region of width K. For narrow learning windows ($K \ll 1/\omega$), this region is small compared to a theta oscillation cycle and much smaller than the width σ of a firing field. Because the envelope of the cross-correlation $C_{ij}(t)$ is a Gaussian with mean T_{ij} and width $\sqrt{2}\sigma$, the slope $C'_{ij}(t = 0)$ scales with the Gaussian factor $G_{T_{ij},\sqrt{2}\sigma}(0) \propto \exp\left[-T_{ij}^2/(4\sigma^2)\right]$. The weight change Δw_{ij} therefore strongly depends on the separation T_{ij} of the firing fields. When the two firing fields do not overlap ($T_{ij} \gg \sigma$), the factor $\exp\left[-T_{ij}^2/(4\sigma^2)\right]$ quickly tends to zero, and sequence learning is not possible. On the other hand, when the two firing fields do have considerable overlap ($T_{ij} \lesssim \sigma$) we have $\exp\left[-T_{ij}^2/(4\sigma^2)\right] \lesssim 1$. In this case, sequence learning may be feasible with narrow learning windows. In this section, we will proceed the mathematical analysis for overlapping fields, which allows us to assume $\exp\left[-T_{ij}^2/(4\sigma^2)\right] \approx 1$.

For overlapping firing fields ($T_{ij} \lesssim \sigma$), let us now consider the fine structure of the cross-correlation $C_{ij}(t)$ for $|t| < K$, as illustrated in Fig. 2D. Importantly, phase precession causes the first positive peak (i.e., for $t > 0$) of C_{ij} to occur at time $c\,T_{ij}$ with $c \ll 1$ (31, 32); phase precession also increases the slope $C'_{ij}(t)$ around $t = 0$, which could be beneficial for temporal-order learning according to Equation 3. To quantify this effect, we calculated the cross-correlation's slope at $t = 0$ (see also Equation 18 in *Materials and Methods*):

$$C'_{ij}(0) \propto G_{T_{ij},\sqrt{2}\sigma}(0) \left[\frac{T_{ij}}{\sigma} + \omega\sigma \sin\left(\omega c T_{ij}\right) + \frac{T_{ij}}{2\sigma}\cos(\omega c T_{ij})\right]. \tag{4}$$

How does $C'_{ij}(0)$ depend on the temporal separation T_{ij} of the firing fields? If the two fields overlap entirely ($T_{ij} = 0$) the sequence has no defined temporal order, and thus $C'_{ij}(0)$ is zero. For at least partly overlapping firing fields ($T_{ij} \lesssim \sigma$) and typical phase precession where $c = \pi/(4\omega\sigma) \ll 1$, we will show in the next paragraph (and explain in *Materials and Methods* in the text below Equation 18) that the second addend in Equation 4 dominates the other two. In this case, $C'_{ij}(0)$ is much higher as compared to the cross-correlation slope in the absence of phase precession ($c = 0$), leading to a clearly larger synaptic weight change for phase precession. The

maximum of $C'_{ij}(0)$ is mainly determined by this second addend (multiplied by $G_{T_{ij},\sqrt{2}\sigma}(0)$) and it can be shown (see *Materials and Methods*) that this maximum is located near $T_{ij} \approx \sqrt{2}\sigma$.

The increase of $C'_{ij}(0)$ induced by phase precession can be exploited by learning windows W that are narrower than a theta cycle (e.g. gray regions in Fig. 2C,D,E). To quantify this effect, let us consider a simple but generic shape of a learning window, for example, the odd STDP window $W(t) = \mu\,\mathrm{sign}(t)\exp(-|t|/\tau)$ with time constant τ and learning rate $\mu > 0$ (Fig. 2E); this STDP window is narrow for $\tau \ll 1/\omega$. Equations 3 and 4 then lead to (see *Materials and Methods*, Equation 19) the average weight change

$$\Delta w_{ij} = A^2 \mu\, \tau^2 \, \frac{G_{T_{ij},\sqrt{2}\sigma}(0)}{\sigma} \left[\frac{T_{ij}}{\sigma} + \frac{\omega\sigma \sin(\omega c T_{ij})}{1+\omega^2\tau^2} + \frac{T_{ij}}{2\sigma}\cos(\omega c T_{ij}) \cdot \frac{1-\omega^2\tau^2}{(1+\omega^2\tau^2)^2} \right], \tag{5}$$

where A depicts the number of spikes per field traversal. Note that, according to Equation 3, the weight change Δw_{ij} in Equation 5 can be interpreted as a time-averaged version of $C'_{ij}(t)$ near $t = 0$ from Equation 4. Thus, Equations 4 and 5 have a similar structure, but Equation 5 includes multiple incidences of the term $\omega^2\tau^2$ that account for this averaging. This term is small for narrow learning windows ($\tau \ll 1/\omega$) and can thus be neglected ($\omega^2\tau^2 \ll 1$) in this limiting case; however, for typical biological values of $\tau \geq 10$ ms and $\omega = 2\pi \cdot 10$ Hz, the peculiar structure of the $\omega^2\tau^2$-containing factor in the third addend in the square brackets is the reason why this addend can be neglected compared to the first one; as a result, the cases of "phase locking" ($c = 0$) and "no theta" (only the first addend remains) are basically indistinguishable. Moreover, for narrow odd learning windows, Δw_{ij} in Equation 5 inherits a number of properties from $C'_{ij}(0)$ in Equation 4: the second addend still remains the dominant one for $T_{ij} \lesssim \sigma$; inherited are also the absence of a weight change for fully overlapping fields ($\Delta w_{ij} = 0$ for $T_{ij} = 0$), the maximum weight change for $T_{ij} \approx \sqrt{2}\sigma$, and $\Delta w_{ij} \to 0$ for $T_{ij} \to \infty$ (Fig. 3A). Furthermore, the prefactor $A^2\mu\tau^2$ in Equation 5 suggests that the average weight change increases with increasing width τ of the learning window, but we emphasize that this increase is restricted to $\tau \ll 1/\omega$ (as we assumed for the derivation), which prohibits a generalization of the quadratic scaling to large τ; the exact dependence on τ will be explained later.

To quantify how much better a sequence can be learned with phase precession as compared to phase locking, we use the ratio of the weight change Δw_{ij} with phase precession ($c > 0$) and the weight change $\Delta w_{ij}(c = 0)$

without phase precession (Fig. 3A), and define the benefit B of phase precession as

$$B := \frac{\Delta w_{ij}}{\Delta w_{ij}(c=0)} - 1. \tag{6}$$

By inserting Equation 5 in Equation 6, we can explicitly calculate the benefit B of phase precession (see Equation 20 in *Materials and Methods* and solid line in Fig. 3B). For $T_{ij} \lesssim \sigma$ and $\omega^4\tau^4 \ll 1$ (see *Materials and Methods*) the benefit B is well approximated by a Taylor expansion up to third order in T_{ij} (dashed line in Fig. 3B),

$$B \approx \frac{2}{3}\omega^2\sigma^2 c \left[1 - \left(\frac{c}{4\sigma^2} \cdot \frac{1-\omega^2\tau^2}{1+\omega^2\tau^2} + \frac{1}{6}\omega^2 c^2\right) T_{ij}^2\right] . \tag{7}$$

The maximum of B as a function of T_{ij} is obtained for $T_{ij} = 0$ (fully overlapping fields), but the average weight change Δw_{ij} is zero at this point. We note, however, that B decays slowly with increasing T_{ij}, so $B\,(T_{ij} = 0)$ can be used to approximate the benefit for small field separations T_{ij} (i.e. largely overlapping fields). For narrow ($\omega\tau \ll 1$) odd STDP windows and slope-size matching ($\omega\sigma c = \pi/4$), we find the maximum $B_{\max} \approx \omega\sigma/2$, which has an interesting interpretation: If we relate σ to the field size L of a Gaussian firing field through $L = 4\sigma$ and if we relate the frequency ω to the period T_θ of a theta oscillation cycle through $T_\theta = 2\pi/\omega$, we obtain $B_{\max} \approx 0.82\, L/T_\theta$, i.e., the maximum benefit of phase precession is about the number of theta oscillation cycles in a firing field. The example in Figure 3B (with firing fields in Fig. 2A) has the maximum benefit $B_{\max} \approx 10$ and the benefit remains in this range for partly overlapping firing fields ($0 < T_{ij} \lesssim \sigma$). We thus conclude that phase precession can boost temporal-order learning by about an order of magnitude for typical cases in which learning windows are narrower than a theta oscillation cycle and overlapping firing fields are an order of magnitude wider than a theta oscillation cycle.

So far, we have considered "average" weight changes that resulted from neural activity that was described by a deterministic firing rate. However, neural activity often shows large variability, i.e., different traversals of the same firing field typically lead to very different spike trains. To account for such variability, we have simulated neural activity as inhomogeneous Poisson processes (see *Materials and Methods* for details). As a result, the change of the weight of a synapse, which depends on the correlation between spikes of the presynaptic and the postsynaptic cells, is a stochastic variable. It is important to consider the variability of the weight change ("noise") in order to assess the significance of the average weight change. For this reason, we utilize the signal-

to-noise ratio (SNR), i.e., the mean weight change divided by its standard deviation (see *Materials and Methods* for details). To do so, we perform stochastic simulations of spiking neurons and calculate the average weight change and its variability across trials. This is done for phase-precessing as well as phase-locked activity. To connect this approach to our previous results, we confirm that the average weight changes estimated from many fields traversals matches well the analytical predictions (Fig. 3A and B, see *Materials and Methods* for details).

The SNR shown in Figure 3C summarizes how reliable is the learning signal in a single traversal of the two firing fields — for the assumed odd learning window. The SNR further depends on T_{ij} and follows a similar shape as the weight changes in Figure 3A. For phase precession, there is a maximum SNR that is slightly shifted to larger T_{ij}; for phase locking, SNR is always much lower. For the synapse connecting two cells with firing fields as in Figure 2A where $T_{ij} = \sigma$, we find a SNR of 0.27, which is insufficient for a reliable representation of a sequence.

To allow reliable temporal-order learning, one possible solution is to increase number of spikes per field traversal A (SNR $\propto \sqrt{A}$, as shown in the Supplemental Information). Another possibility is to increase the number of synapses. In *Materials and Methods* we show that SNR $\propto \sqrt{M}$ where M is the number of identical and uncorrelated synapses. Therefore, to achieve SNR $\gtrsim 1$ for $A = 10$, one needs $M \gtrsim 14$ synapses.

In summary, for narrow, odd learning windows ($\tau \ll 1/\omega \ll \sigma$), temporal-order learning could benefit tremendously from phase precession as long as firing fields have some overlap. Average weight changes and the SNR are highest, however, for clearly distinct but still overlapping firing fields. It should be noted that any even component of the learning window would increase the noise and thus further decrease the SNR.

Dependence of temporal-order learning on the width of the learning window for overlapping firing fields

To investigate how temporal-order learning for an odd learning window depends on its width, we vary the parameter τ and quantify the average synaptic weight change Δw_{ij} and the SNR both analytically and numerically. We first study overlapping firing fields (Fig. 4) and later consider non-overlapping firing fields (Fig. 5).

For partly overlapping firing fields (e.g., $T_{ij} = \sigma$), we find numerically that the average synaptic weight

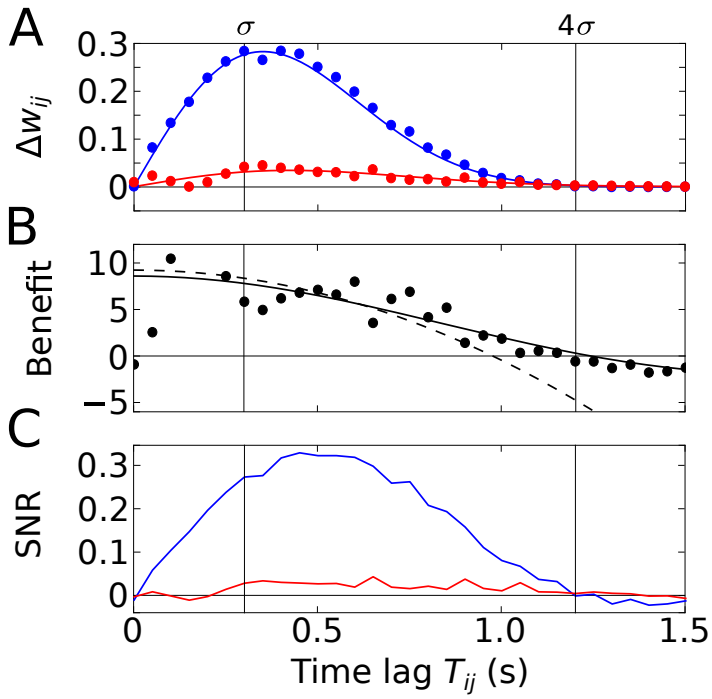

Figure 3: Temporal-order learning for narrow learning windows ($\tau \ll \frac{1}{\omega}$). (A) The average synaptic weight change Δw_{ij} depends on the temporal separation T_{ij} between the firing fields. Phase precession (blue) yields higher weight changes than phase locking (red). Simulation results (circles, averaged across 10^4 repetitions) and analytical results (lines, Equation 5) match well. The vertical lines mark time lags of σ and 4σ, respectively, where 4σ approximates the total field width. **(B)** The benefit B of phase precession is determined by the ratio of the average weight changes of two scenarios from (A). The solid and dashed lines depict the analytical expression for the benefit (Equation 20) and its approximation for small T_{ij} (Equation 7), respectively. **(C)** Signal-to-noise ratio (SNR) of the weight change as a function of the firing-field separation T_{ij}. The SNR is defined as the mean weight change divided by the standard deviation across trials in the simulation. Colors as in (A). Parameters for all plots: $\omega = 2\,\pi \cdot 10$ Hz, $\sigma = 0.3$ s, $\tau = 10$ ms, $\mu = 1$, $c = 0.042$, $A = 10$.

change Δw_{ij} (the "learning signal") increases monotonically for increasing τ and saturates (colored curves in Fig. 4A). This is because for increasing τ the overlap between the learning window and the cross-correlation function grows, and this overlap begins to saturate as soon as the learning window is wider than T_{ij}, i.e., the value at which the cross-correlation assumes its maximum (cmp. dashed curve in Fig. 2C). To analytically calculate the saturation value of Δw_{ij} for large learning-window widths ($\tau \gg \sigma$), we can approximate the learning window

as a step function (see *Materials and Methods* for details) and find the maximum

$$\Delta w_{ij}^{\max} \approx A^2 \mu \, \mathrm{erf}\left(\frac{T_{ij}}{2\sigma}\right) \tag{8}$$

that provides an upper bound to the weight change for overlapping firing fields (solid line in Fig. 4A). For $\tau \lesssim 1/\omega$ (and actually well beyond this region), the analytical small-tau approximation of Δw_{ij} (Equation 5, dashed curves in Fig. 4A) matches the numerical results well.

The results in Figure 4A confirm that Δw_{ij} is increased by phase precession for narrow learning windows but is independent of phase precession for $\tau \gg 1/\omega$. Thus, the benefit B becomes small for large τ (Fig. 4B) because, for large enough τ, the theta oscillation completes multiple cycles within the width of the learning window. To better understand this behavior, let us return to Equation 1: if the product of a *wide* learning window and the cross-correlation C_{ij} is integrated to obtain the weight change, the oscillatory modulation of the cross-correlation (e.g., as in Fig. 2C) become irrelevant; similarly, according to Equation 3, the particular value of the derivative $C'_{ij}(t)$ near $t = 0$ can be neglected. Consequently, for $\tau \gg 1/\omega$ phase precession and phase locking as well as the scenario of firing fields that are not theta modulated yield the same weight change (Fig. 4A), and the benefit approaches 0 (Fig. 4B). Wide learning windows thus ignore the temporal (theta) fine-structure of the cross-correlation.

How noisy is this learning signal Δw_{ij} across trials? Figure 4C shows that for odd learning windows the SNR increases with increasing τ and, for $\tau \gg \frac{1}{\omega}$, approaches a constant value. This constant value is the same for phase precession, phase locking, or no theta oscillations at all. Taken together, for large enough τ, the advantage of phase precession vanishes. For small enough τ, phase precession increases the SNR, which confirms and generalizes the results in Fig. 3C. Remarkably, the SNR for "phase locking" is lower than the one for "no theta", which means that theta oscillations without phase precession degrade temporal-order learning, even though theta oscillations as such were emphasized to improve the modification of synaptic strength in many other cases (e.g. refs. 9, 33).

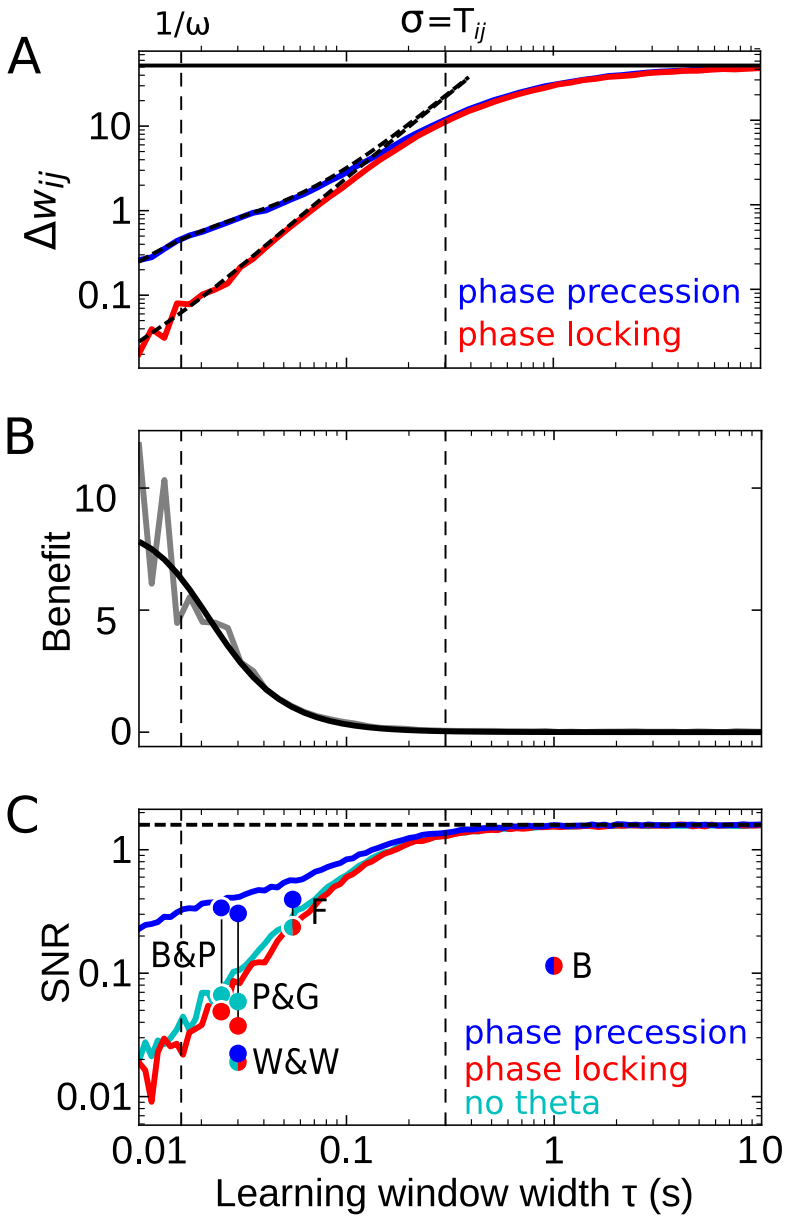

Figure 4: Effect of the learning-window width on temporal-order learning for overlapping fields (here: $T_{ij} = \sigma$). [Caption on the next page.]

Figure 4: Effect of the learning-window width on temporal-order learning for overlapping fields (here: $T_{ij} = \sigma$). **(A)** Average weight change Δw_{ij} as a function of width τ (for the asymmetric window W in Equation 14) for phase precession and phase locking (colored curves). The solid black line depicts the theoretical maximum for large τ ($\Delta w_{ij} \approx 52$, Equation 8). The dashed curves show the analytical small-tau approximations (Equation 5). The vertical dashed lines mark $1/\omega \approx 0.016$ s and the value of $\sigma = T_{ij} = 0.3\ s$, respectively. **(B)** The benefit B of phase precession is largest for narrow learning windows, and it approaches 0 for wide windows. Simulations (gray line) and analytical result (black line, small-tau approximation from Equation 20) match well. **(C)** The signal-to-noise ratio (SNR; phase precession: blue, phase locking: red, no theta: cyan) takes into account that only the asymmetric part of the learning window is helpful for temporal-order learning. For large τ, all three coding scenarios induce the same SNR. The horizontal dashed black line depicts the analytical limit of the SNR for large τ and overlapping firing fields (SNR ≈ 1.6, Equation 17 of the Supplemental Information). Dots represent the SNR for experimentally observed learning windows. The learning windows were taken from Bi and Poo ("B&P", ref. 15: their Fig. 1), Froemke et al. ("F", ref. 16: their Fig. 1D bottom), Wittenberg and Wang ("W&W", ref. 17: their Fig. 3E), Pfister and Gerstner ("P&G", ref. 34: their Table 4, 'All to All', 'minimal model'), and Bittner et al. ("B", ref. 35: their Fig. 3D). For "B&P", "F", and "B", the position of the dots on the horizontal axis was estimated as the average time constants for positive and negative lobes of the learning windows. Wittenberg and Wang modelled their learning rule by a difference of Gaussians — we approximated the corresponding time constant as 30 ms. For the triplet rule by Pfister and Gerstner, we used the average of three time constants: the two pairwise-interaction time constants (as in Bi and Poo) and the triplet-potentiation time constant. Parameters for all plots: $T_{ij} = 0.3$ s, $\omega = 2\ \pi \cdot 10$ Hz, $\sigma = 0.3$ s, $c = 0.042$, $A = 10$, $\mu = 1$. Colored/gray curves and dots are obtained from stochastic simulations; see *Materials and Methods* for details.

Figure 4C predicts that a large τ yields the biggest SNR, and thus wide learning windows are the best choice for temporal-order learning; however, we note that this conclusion is restricted to odd (i.e., asymmetric) learning windows. An additional even (i.e., symmetric) component of a learning window would increase the

noise without affecting the signal, and thus would decrease the SNR (dots in Fig. 4C). It is remarkable that the only experimentally observed instance of a wide window (with $\tau \approx 1$s in ref. 35) has a strong symmetric component, which leads to a low SNR (dot marked "B" in Fig. 4C).

Taken together, we predict that temporal-order learning would strongly benefit from wide, asymmetric windows. However, to date, all experimentally observed (predominantly) asymmetric windows are narrow (e.g. refs. 15–17; see refs. 15, 18 for reviews).

Temporal-order learning for wide learning windows ($K \gg \sigma$)

We finally restrict our analysis to wide learning windows, which allows us then to also consider non-overlapping firing fields (Figure 5A,we again use two Gaussians with widths σ and separation T_{ij}). To allow for temporal-order learning in this case, the spikes of two non-overlapping fields can only be "paired" by a wide enough learning window. As already indicated in Fig. 4, phase precession does not affect the weight change for such wide learning windows where the width τ of the learning window obeys $\tau \gg 1/\omega$ (note that we always assumed many theta oscillation cycles within a firing field, i.e., $1/\omega \gg \sigma$). Furthermore, Fig. 4 indicated that only the asymmetric part of the learning window contributes to temporal-order learning. For the analysis of temporal-order learning with non-overlapping firing fields and wide learning windows, we thus ignore any theta modulation and phase precession and evaluate, again, only the odd STDP window $W(t) = \mu\,\mathrm{sign}(t)\exp(-|t|/\tau)$. In this case, the weight change (Equation 1) is still determined by the cross-correlation function and the learning window (examples in Fig. 5B, C). The resulting weight change Δw_{ij} as a function of the temporal separation T_{ij} of firing fields is shown in Fig. 5D: with increasing T_{ij}, the weight Δw_{ij} quickly increases, reaches a maximum, and slowly decreases. The initial increase is due to the increasing overlap of the Gaussian bump in C_{ij} with the positive lobe of the learning window. The decrease, on the other hand, is dictated by the time course of the learning window. For $\tau \gg \sigma$, these two effects can be approximated by

$$\Delta w_{ij} \approx A^2 \mu\, \mathrm{erf}\left(\frac{T_{ij}}{2\sigma}\right) \exp\left(-\frac{T_{ij}}{\tau}\right) \tag{9}$$

in which the error function describes the overlap of the cross-correlation with the learning window and the exponential term describes the decay of the learning window (dashed black curve in Fig. 5D, see also Equation 25

in *Materials and Methods* for details).

How does the SNR of the weight change depend on the separation T_{ij} of firing fields? For $T_{ij} = 0$, the signal is zero and thus also the SNR. As T_{ij} increases, both signal and noise increase, but quickly settle on a constant ratio. The value of the SNR height of this plateau can be approximated by

$$\mathrm{SNR} \approx \frac{A}{\sqrt{2A+1}} \tag{10}$$

(dashed line in Fig. 5E), where A is the number of spikes within a firing field (Equation 11). For $A = 10$, we find SNR ≈ 2.2, allowing for temporal-order learning with a single synapse. We note that this conclusion is limited to asymmetric STDP windows. A symmetric component (like in ref. 35) decreases the SNR and makes temporal-order learning less efficient.

Taken together, temporal-order learning can be performed with wide STDP windows, and phase precession does not provide any benefit; but temporal-order learning requires a purely asymmetric plasticity window. For non-overlapping firing fields, wide learning windows are essential to bridge a temporal gap between the fields.

Discussion

In this report, we show that phase precession facilitates the learning of the temporal order of behavioral sequences for asymmetric learning windows that are shorter than a theta cycle. To quantify this improvement, we use additive, pairwise STDP and calculate the expected weight change for synapses between two activated cells in a sequence. We confirm the long-held hypothesis (13) that phase precession bridges the vastly different time scales of the slow sequence of behavioral events and the fast STDP rule. Synaptic weight changes can be an order of magnitude higher when phase precession organizes the spiking of multiple cells at the theta time scale as compared to phase-locking cells.

Other mechanisms and models for sequence learning

As an alternative mechanism to bridge the time scales of behavioral events and the induction of synaptic plasticity, Drew and Abbott (4) suggested STDP and persistent activity of neurons that code for such events. The authors assume regularly firing neurons that slowly decrease their firing rate after the event and show that

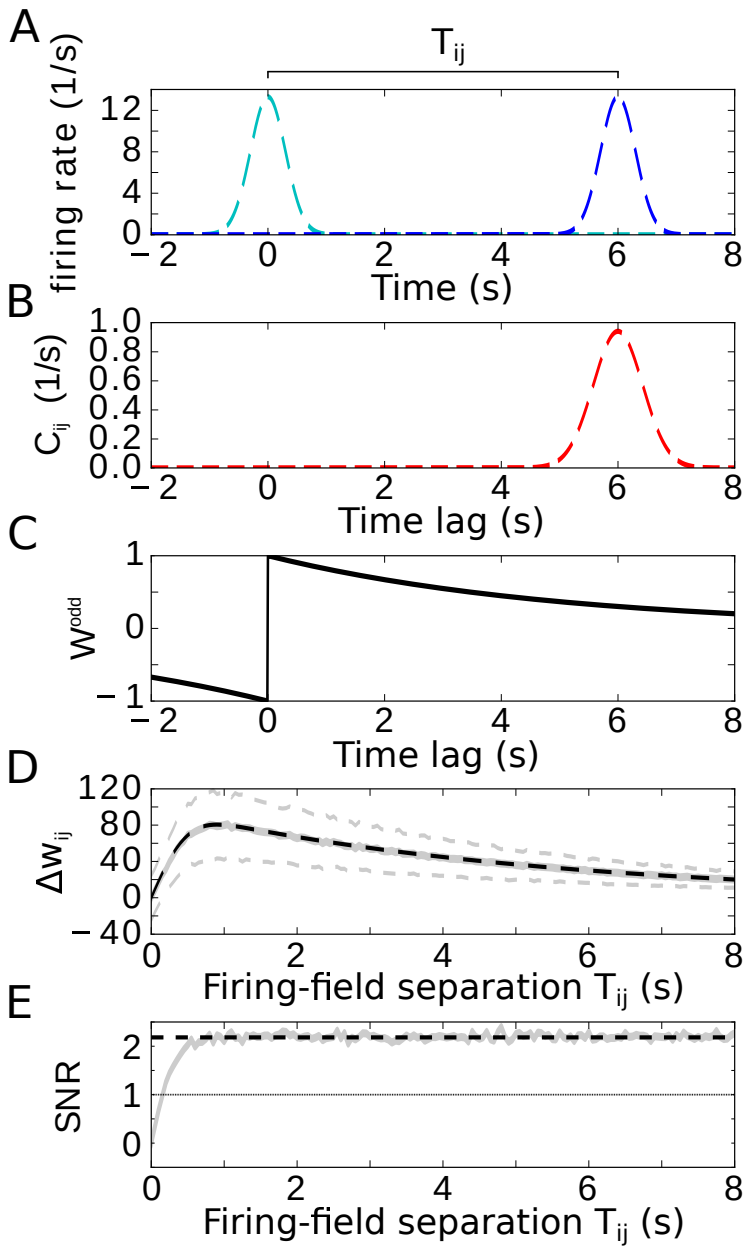

Figure 5: Temporal-order learning for non-overlapping firing fields using wide, asymmetric learning windows. **(A)** Firing rates of two example cells with non-overlapping firing fields. **(B)** Cross-correlation C_{ij} of the two cells from (A). **(C)** Asymmetric learning window with large width ($\tau = 5$s). **(D)** Resulting weight change Δw_{ij} for wide learning window and non-overlapping firing fields. The solid grey line depicts the average weight change. The dashed grey lines represent ± 1 standard deviation across 1000 repetitions of stochastic spiking simulations. The analytical curve (dashed black line, Equation 9) matches the simulation results. **(E)** SNR of the weight change. Results of the stochastic simulations are shown by the grey curve. The SNR saturates for larger T_{ij}, which fits the analytical expectation (dashed black line, Equation 10). Parameters, unless varied in a plot: $T_{ij} = 6$s, $\sigma = 0.3$ s, $\tau = 5$ s, $\mu = 1$, $A = 10$.

this leads to a temporal compression of the sequence of behavioral events. For stochastically firing neurons, this approach is similar to ours with two overlapping, unmodulated Gaussian firing fields. In this case, sequence learning is possible, but the efficiency can be improved considerably by phase precession.

Sato and Yamaguchi (20) as well as (22) investigated the memory storage of behavioral sequences using phase precession and STDP in a network model. In computer simulations they find that phase precession facilitates sequence learning, which is in line with our results. In contrast to these approaches, our study focuses on a minimal network (two cells), but this simplification allows us to (i) consider a biologically plausible implementation of STDP, firing fields, and phase precession and (ii) derive analytical results. These mathematical results predict parameter dependencies, which is difficult to achieve with only computer simulations.

Related to our work is also the approach by Masquelier and colleagues (36) who showed that pattern detection can be performed by single neurons using STDP and phase coding, yet they did not include phase precession. They consider patterns in the input whereas, in our framework, it might be argued that patterns between input and output are detected instead.

Noisy activity of neurons and prediction of the minimum number of synapses for temporal-order learning

To account for stochastic spiking, we use Poisson neurons. We find that a single synapse is not sufficient to reliably encode a minimal two-neuron sequence in a single trial because the fluctuations of the weight change are too large. Fortunately, the SNR scales with $\sqrt{MA}$, i.e., the square root of the number M of identical, but independent synapses and the number A of spikes per field traversal of the neurons. For generic hippocampal place fields and typical STDP, we predict that about 14 synapses are sufficient to reliably encode temporal order in a single traversal. Interestingly, peak firing rates of place fields are remarkably high (up to 50 spikes/s; e.g. 8, 37). Taken together, in hippocampal networks, reliable encoding of the temporal order of a sequence is possible with a low number of synapses, which matches simulation results on memory replay (24).

Width, shape, and symmetry of the STDP window are critical for temporal-order learning

Various widths have been observed for STDP learning windows (15, 18). We show that for all experimentally found STDP time constants phase precession can improve temporal-order learning. However, for learning windows much wider than a theta oscillation cycle, the benefit of phase precession for temporal-order learning is small. Wide learning windows, where the width can be even on a behavioral time scale of $\approx$ 1s (35) or larger, could, on the other hand, enable the association of non-overlapping firing fields. Alternatively, non-overlapping firing fields might also be associated by narrow learning windows if additional cells (with firing fields that fill the temporal gap) help to bridge a large temporal difference, much like "time cells" in the hippocampal formation (reviewed in 38).

STDP windows typically have symmetric and asymmetric components (18). We find that only the asymmetric component supports the learning of temporal order. In contrast, the symmetric component strengthens both forward and backward synapses by the same amount and thus contributes to the association of behavioral events independent of their temporal order. For example, the learning window reported by Bittner et al. (35) shows only a mild asymmetry and is thus unfavorable to store the temporal order of behavioral events. Only long, predominantly asymmetric STDP windows would allow for effective temporal-order learning (cf. Fig. 4).

Generally, the shape of STDP windows is subject to neuromodulation; for example, cholinergic and adrenergic modulation can alter its polarity and symmetry (39). Also dopamine can change the symmetry of the learning window (40). Therefore, sequence learning could be modulated by the behavioral state (attention, reward, etc.) of the animal.

Key features of phase precession for temporal order-learning: generalization to non-periodic modulation of activity

For STDP windows narrower ($\lesssim$ 10 ms) than a theta cycle ($\gtrsim$ 100 ms), we argue that the slope of the cross-correlation function at zero offset controls the change of the weight of the synapse connecting two neurons; and we show that phase precession can substantially increase this slope. This result predicts that features of the cross-correlation at temporal offsets that are larger than the width of the learning window are irrelevant for

temporal-order learning. It is thus conceivable to boost temporal-order learning even without phase precession, which is weak if theta oscillations are weak, as for example in bats (41) and humans (42). In this case, temporal-order learning may instead benefit from two other phenomena that could create an appropriate shape of the cross-correlation: (1) Spiking of cells is locked to common (aperiodic) fluctuations of excitability. (2) Each cell responds the faster to an increase in its excitability the longer ago its firing field has been entered, which may be mediated by a progressive facilitation mechanism. Together, these phenomena can make the cross-correlation exhibit a steeper slope around zero and could even give rise to a local maximum at a positive offset. This temporal fine structure is superimposed on a slower modulation, which is related to the widths of the firing fields. In summary, a progressively decreasing delay of spiking with respect to non-rhythmic fluctuations in excitation generalizes the notion of phase precession. Interestingly, synaptic short-term facilitation, which could generate the described fine structure of the cross-correlation, has also been proposed as mechanism underlying phase precession (43).

Model assumptions

In our model we assumed that recurrent synapses (e.g. between neurons representing a sequence) are plastic but weak during encoding, such that they have a negligible influence on the postsynaptic firing rate; and that the feedforward input dominates neuronal activity. These assumptions seem justified as Hasselmo (39) indicated that excitatory feedback connections may be suppressed during encoding to avoid interference from previously stored information (see also 44). Furthermore, neuromodulators facilitate long-term plasticity (reviewed, e.g., by 45), which also supports our assumptions.

The assumption of weak recurrent connections implies that these connections do not affect the dynamics. Consequently (and in contrast to ref. 14), we thus hypothesize that phase precession is not generated by the local, recurrent network (see also, e.g. 46); instead, we assume that phase precession is inherited from upstream feedforward inputs (47, 48) or generated locally by a cellular/synaptic mechanism (49–52). After temporal-order learning was successful, the resulting asymmetric connections could indeed also generate phase precession (as demonstrated by the simulations in ref. 14), and this phase precession could then even be similar to the one that

has initially helped to shape synaptic connections. Finally, inherited or local cellularly/synaptically-generated phase precession and locally network-generated phase precession could interact (as reviewed, for example in ref. 53).

We assumed in our model that the widths of the two firing fields that represent two events in a sequence are identical (see, e.g., Fig. 2A). But firing fields may have different widths, and in this case a slope-size matched phase precession would fail to reproduce the timing of spikes required for the learning of the correct temporal order of the two events. For example, the learned temporal order of events (timed according to field entry) would even be reversed if two fields with different sizes are aligned at their ends. How could the correct temporal order nevertheless be learned in our framework? In the hippocampus, theta oscillations are a traveling wave (54, 55) such that there is a positive phase offset of theta oscillations for the wider firing fields in the more ventral parts of the hippocampus. This traveling-wave phenomenon could preserve the temporal order in the phase-precession-induced compressed spike timing, as also pointed out earlier (56, 57).

Our results on learning rules for sequence learning rely on pairwise STDP in which pairs of presynaptic and postsynaptic spikes are considered. Conversely, triplet STDP considers also motifs of three spikes (either 2 presynaptic - 1 postsynaptic or 2 postsynaptic - 1 presynaptic) (34). Triplets STDP models can reproduce a number of experimental findings that pairwise STDP could not, for example the dependence on the repetition frequency of spike pairs (58). To investigate the influence of triplet interactions on sequence learning, we implemented the generic triplet rule by Pfister and Gerstner (34). We used their "minimal" model, which was regarded as the best model in terms of number of free parameters and fitting error; for the parameters they obtained from fitting the triplet STDP model to hippocampal data, we found only mild differences to our results (see, e.g., Fig. 4C). Differences are small because the fitted time constant of the triplet term (40 ms) is smaller than typical inter-spike intervals ($\gtrsim$ 50 ms, minimum in field centers) in our simulations.

Replay of sequences and storage of multiple and overlapping sequences

A sequence imprinted in recurrent synaptic weights can be replayed during rest or sleep (27, 59–62), which was also observed in network-simulation studies (25, 26, 63). Replay could thus be a possible readout of the

temporal-order learning mechanism. However, replay depends on the many parameters of the network, and a thorough investigation of is beyond the scope of this manuscript. Therefore, we focus on synaptic weight changes that represent the formation of sequences in the network, which underlies replay, and we do not simulate replay.

We have considered the minimal example of a sequence of two neurons. Sequences can contain many more neurons, and the question arises how two different sequences can be told apart if they both contain a certain neuron, but proceed in different directions — as they might do for sequences of spatial or non-spatial events (64). In this case, it may be beneficial to not only strengthen synapses that connect direct successors in the sequence but also synapses that connect the second-to-next neuron. In this way, the two crossing sequences could be disambiguated, and the wider context in which an event is embedded becomes associated, which is in line with retrieved-context theories of serial-order memory (65). More generally, it is an interesting question of how many sequences can be stored in a network of a given size. Gillett et al. (26) were able to analytically calculate the storage capacity for the storage of sequences in a Hebbian network.

In conclusion, our model predicts that phase precession enables efficient and robust temporal-order learning. To test this hypothesis, we suggest experiments that modulate the shape of the STDP window or selectively manipulate phase precession and evaluate memory of temporal order.

Materials and Methods

Experimental Design: Model description

We model the time-dependent firing rate of a phase precessing cell i (two examples in Fig. 2A) as

$$f_i(t) = A \cdot G_{\mu_i,\sigma}(t) \cdot \{1 + \cos[\omega(t - c\mu_i)]\} \ , \tag{11}$$

where the scaling factor A determines the number of spikes per field traversal and $G_{\mu_i,\sigma}(t) = 1/(\sqrt{2\pi}\sigma)\cdot\exp[-(t-\mu_i)^2/(2\sigma^2)]$ is a Gaussian function that describes a firing field with center at μ_i and width σ. The firing field is sinusoidally modulated with theta frequency ω (but the sinusoidal modulation is not a critical assumption, see *Discussion*), with typically many oscillation cycles in a firing field ($\omega\sigma \gg 1$). The compression factor c can be used to vary between phase precession ($c > 0$), phase locking ($c = 0$), and phase recession ($c < 0$) because

the average population activity of many such cells oscillates at frequency of $(1-c)\omega$ (31, 33), which provides a reference frame to assign theta phases (Fig. 2A). Usually, $|c| \ll 1$ with typical values $c \lesssim 1/(\sigma\omega)$ (31); for a pair of cells with overlapping firing fields (centers separated by $T_{ij} := \mu_j - \mu_i$) the phase delay is $\omega c T_{ij}$ (Fig. 2B).

To quantify temporal-order learning, we consider the average weight change Δw_{ij} of the synapse from cell i to cell j, which is (30)

$$\Delta w_{ij} = \int_{-\infty}^{\infty} \mathrm{d}t\, W(t)\, C_{ij}(t) \tag{12}$$

where $C_{ij}(t)$ is the cross-correlation between the firing rates f_i and f_j of cells i and j, respectively (Fig. 2C,D):

$$C_{ij}(t) = \int_{-\infty}^{\infty} \mathrm{d}t'\, f_i(t') f_j(t+t')\,. \tag{13}$$

$W(t)$ denotes the synaptic learning window, for example the asymmetric window

$$W(t) = \mu \begin{cases} +\exp(-t/\tau), & t \geq 0 \\ -\exp(+t/\tau), & t < 0\,, \end{cases} \tag{14}$$

where τ is the time constant and $\mu > 0$ is the learning rate (Fig. 2E).

For the following calculations, we make two assumptions that are reasonable in the hippocampal formation (8, 15, 31) :

(i) The theta oscillation has multiple cycles within the Gaussian envelope of the firing field in Equation 11 $(1/\omega \ll \sigma)$.

(ii) The window W is short compared to the theta period $(\tau \ll 1/\omega)$.

Analytical approximation of the cross-correlation function

To explicitly calculate the cross-correlation $C_{ij}(t)$ as defined in Equation 13, we plug in the firing-rate functions (Equation 11) for the two neurons:

$$
\begin{aligned}
C_{ij}(t) = \int_{-\infty}^{\infty} \mathrm{d}t'\, A \cdot G_{0,\sigma}(t') \cdot [1 + \cos(\omega t')] \\
\cdot A \cdot G_{T_{ij},\sigma}(t+t') \cdot \left\{1 + \cos\left[\omega(t + t' - cT_{ij})\right]\right\} \\
= A^2 \int_{-\infty}^{\infty} \mathrm{d}t' \Bigg\{ & G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \\
+ \quad & G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \cos(\omega t') \\
+ \quad & G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \cos\left[\omega(t + t' - cT_{ij})\right] \\
+ \quad & G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \cos(\omega t') \cos\left[\omega(t + t' - cT_{ij})\right] \Bigg\} .
\end{aligned}
$$

The first term (out of four) describes the cross-correlation of two Gaussians, which results in a Gaussian function centered at T_{ij} and with width $\sigma\sqrt{2}$. For the second term, we note that the product of two Gaussians yields a function proportional to a Gaussian with width $\sigma/\sqrt{2}$, and then use assumption (i). When integrated, the second term's contribution to $C_{ij}(t)$ is negligible because the cosine function oscillates multiple times within the Gaussian bump, i.e., positive and negative contributions to the integral approximately cancel. The same argument applies to the third term. For the fourth term, we use the trigonometric property $\cos(\alpha) \cdot \cos(\beta) = \frac{1}{2}\left(\cos(\alpha+\beta) + \cos(\alpha-\beta)\right)$. We set $\alpha = \omega t'$, $\beta = \omega(t + t' - cT_{ij})$ and find

$$
\begin{aligned}
& G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \cos(\omega t') \cos[\omega(t + t' - cT_{ij})] = \\
& \frac{1}{2} G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \cos[\omega(t + 2t' - cT_{ij})] \\
+ & \frac{1}{2} G_{0,\sigma}(t') G_{T_{ij},\sigma}(t+t') \cos[\omega(t - cT_{ij})] .
\end{aligned}
$$

Again, we use assumption (i) and neglect the first addend on the right-hand side. Notably, the cosine function in the second addend is independent of the integration variable t'. Taken together, we find

$$
C_{ij}(t) \approx \; A^2\, G_{T_{ij},\sigma\sqrt{2}}(t) \left\{1 + \frac{1}{2} \cos[\omega(t - cT_{ij})]\right\} . \tag{15}
$$

Thus, the cross-correlation can be approximated by a Gaussian function (center at T_{ij}, width $\sigma\sqrt{2}$) that is theta modulated with an amplitude scaled by the factor $\frac{1}{2}$.

To further simplify Equation 15, we note that the time constant τ of the STDP window is usually small compared to the theta period (assumption (ii), Fig. 2C,D,E). Structures in $C_{ij}(t)$ for $|t| \gg \tau$ thus have a negligible effect on the synaptic weight change. Therefore, we can focus on the cross-correlation for small temporal lags. In this range, we approximate the (slow) Gaussian modulation of $C_{ij}(t)$ (Fig. 2C,D, dashed red line) by a linear function, that is,

$$\begin{aligned} G_{T_{ij},\sqrt{2}\sigma}(t) &\approx \frac{\mathrm{d}}{\mathrm{d}t} G_{T_{ij},\sqrt{2}\sigma}(t)\Big|_{t=0} \cdot t + G_{T_{ij},\sqrt{2}\sigma}(0) \\ &= G_{T_{ij},\sqrt{2}\sigma}(0) \left(\frac{T_{ij}}{2\sigma^2} \cdot t + 1 \right) . \end{aligned} \tag{16}$$

Inserting this result in Equation 15, we approximate the cross-correlation function $C_{ij}(t)$ for $|t| \lesssim \tau$ as (Fig. 2D, dashed black line)

$$C_{ij}(t) \approx A^2\, G_{T_{ij},\sqrt{2}\sigma}(0) \left(\frac{T_{ij}}{2\sigma^2} \cdot t + 1 \right) \left[1 + \frac{1}{2} \cos\left(\omega(t - cT_{ij})\right) \right] . \tag{17}$$

In the *Results* we show that the slope of the cross-correlation function at $t = 0$ is important for temporal-order learning. From Equation 17 we find

$$C'_{ij}(0) \approx \frac{A^2 T_{ij}}{2\sigma^2} G_{T_{ij},\sqrt{2}\sigma}(0) \left[1 + \omega\sigma\, \omega c\sigma \frac{\sin\left(\omega c T_{ij}\right)}{\omega c T_{ij}} + \frac{\cos(\omega c T_{ij})}{2} \right] , \tag{18}$$

which has three addends within the square brackets. Let us estimate the relative size of the second and third terms with respect to the first one. The third term is at most of the order of 0.5 because $|\cos(\omega c T_{ij})| \leq 1$. For the second addend, we note that $\sin(\omega c T_{ij})/(\omega c T_{ij})$ approaches 1 for $T_{ij} \to 0$ and remains in this range for $|\omega c T_{ij}| \lesssim \pi/4$. This condition is fulfilled for $|T_{ij}| \lesssim \sigma$ if we assume slope-size matching of phase precession (31), i.e., $\omega c\sigma \approx \frac{\pi}{4} \approx 0.79$. Then, the size of the second addend is dictated by the factor $\omega\sigma$, which is large according to assumption (i). In other words, for typical phase precession and $|T_{ij}| \lesssim \sigma$, the second addend is much larger than the other two.

To further understand the structure of $C'_{ij}(0)$, which is also shaped by the prefactors in front of the square brackets, we first note that $C'_{ij}(0)$ is zero for fully overlapping firing fields ($T_{ij} \to 0$). On the other hand, for very large field separations ($T_{ij} \gg \sigma$), the Gaussian term G causes $C'_{ij}(0)$ to become zero. The prefactors have a maximum at $|T_{ij}| = \sqrt{2}\sigma$. The maximum's exact location is slightly shifted by the second addend but remains near $\sqrt{2}\sigma$. This peak will be important because it is inherited by the average weight change (Equation 3).

Average weight change

Having approximated the cross-correlation function and its slope at zero (Equations 17,18), we are now ready to calculate the average synaptic weight change (Equation 3) for the assumed STDP window (Equation 14). Standard integration methods yield

$$\Delta w_{ij} = \frac{A^2 \cdot \mu \cdot \tau^2 \cdot T_{ij}}{\sigma^2} G_{T_{ij},\sqrt{2}\sigma}(0) \left[1 + \frac{\omega^2\sigma^2 c}{\omega^2\tau^2+1} \frac{\sin(\omega c T_{ij})}{\omega c T_{ij}} + \frac{(1-\omega^2\tau^2)\cos(\omega c T_{ij})}{2(1+\omega^2\tau^2)^2}\right]. \tag{19}$$

Because Δw_{ij} is a temporal average of $C'_{ij}(t)$ for small t (see interpretation of Equation 3), the weight change's structure resembles the previously discussed structure of $C'_{ij}(0)$. The averaging introduces additional factors proportional to $1 \pm \omega^2\tau^2$, but for $\omega\tau \ll 1$ [assumption (ii)] those have only minor effects on the relative size of the three addends. The second term still dominates. Importantly, $\Delta w_{ij} = 0$ for $T_{ij} = 0$ and the position of the peak at $T_{ij} \lesssim \sqrt{2}\sigma$ is inherited from $C'_{ij}(0)$ (cf. Fig. 3A).

The benefit of phase precession

To quantify the benefit B of phase precession, we consider the expression $\Delta w_{ij}/\Delta w_{ij}(c=0) - 1$, because Δw_{ij} describes the overall weight change (including phase precession), and $\Delta w_{ij}(c=0)$ serves as the baseline weight change due to the temporal separation of the firing fields (without phase precession). We subtract 1 to obtain $B = 0$ when the weight changes are the same with and without phase precession. From Equation 19 we find

$$B = \frac{2}{3}\omega^2\sigma^2 c \cdot \frac{\sin(\omega c T_{ij})}{\omega c T_{ij}} \cdot \frac{1+\omega^2\tau^2}{1+\omega^2\tau^2+\frac{2}{3}\omega^4\tau^4} + \frac{\cos(\omega c T_{ij})-1}{3} \cdot \frac{1-\omega^2\tau^2}{1+\omega^2\tau^2+\frac{2}{3}\omega^4\tau^4}. \tag{20}$$

To better understand the structure of B, we Taylor-expand it in T_{ij} up to the third order and assume $\omega^4\tau^4 \ll 1$ [cf. assumption (ii)]. The result is

$$B \approx \frac{2}{3}\omega^2\sigma^2 c \left[1 - \left(\frac{1}{6}\omega^2 c^2 + \frac{c}{4\sigma^2} \cdot \frac{1-\omega^2\tau^2}{1+\omega^2\tau^2}\right) T_{ij}^2\right]. \tag{21}$$

Thus, B assumes a maximum for $T_{ij} = 0$ and slowly decays for small T_{ij} (Fig. 3B). Using slope-size matching

($\omega\sigma c = \pi/4$), the maximal benefit is

$$B_{\max} \approx \frac{\pi}{6}\omega\sigma = \frac{\pi^2}{12}\frac{L}{T_\theta} \approx 0.82\frac{L}{T_\theta}\,, \tag{22}$$

where $L = 4\sigma$ depicts the total field size and $T_\theta = \frac{2\pi}{\omega}$ is the period of the theta oscillation. Thus, the number of theta cycles per firing field determines the benefit for small separations of the firing fields.

Average weight change for wide learning windows

In this paragraph we relax assumption (ii), i.e., we consider wide asymmetric learning windows W (Equation 14 with $\tau \gg \sigma$). Furthermore, we neglect any theta-oscillatory modulation of the firing fields in Equation 11 and, thus, C_{ij} in Equation 15.

First, for non-overlapping fields ($T_{ij} \gg \sigma$), the learning window can be approximated to be constant near the peak of the Gaussian bump of C_{ij}. We can thus rewrite Equation 1 as

$$\Delta w_{ij} \approx W(T_{ij}) \int_{-\infty}^{\infty} C_{ij}(t)\,\mathrm{d}t = A^2\mu\,\exp\left(-\frac{T_{ij}}{\tau}\right). \tag{23}$$

Second, for overlapping fields ($0 < T_{ij} \lesssim \sigma$), the Gaussian bump of C_{ij} partly lies on the negative lobe of W. We can approximate $W(t) = \mathrm{sign}(t)$, and the average weight change in Equation 1 then reads

$$\Delta w_{ij} \approx A^2\mu\,\mathrm{erf}\left(\frac{T_{ij}}{2\sigma}\right). \tag{24}$$

Combining the two limiting cases in Equations 23 and 24 yields

$$\Delta w_{ij} \approx A^2\mu\,\mathrm{erf}\left(\frac{T_{ij}}{2\sigma}\right)\exp\left(-\frac{T_{ij}}{\tau}\right)\,. \tag{25}$$

Signal-to-noise ratio

To correctly encode the temporal order of behavioral events, the average weight change Δw_{ij} of a forward synapse needs to be larger than the average weight change Δw_{ji} of the corresponding backward synapse. We thus define the signal-to-noise ratio as

$$\mathrm{SNR} = \frac{\Delta w_{ij} - \Delta w_{ji}}{\mathrm{std}\left(\Delta w_{ij}^k\right) + \mathrm{std}\left(\Delta w_{ji}^k\right)}\,,$$

where std() denotes the standard deviation and Δw_{ij}^k, Δw_{ji}^k are the weight changes for trial $k \in [1, N]$, the averages across trials being $\Delta w_{ij} = \langle \Delta w_{ij}^k \rangle_k$ and $\Delta w_{ji} = \langle \Delta w_{ji}^k \rangle_k$. This expression for the SNR "punishes" the non-sequence-specific strengthening of backward synapses. Specifically, SNR = 0 for a symmetric (even) learning window, because the numerator (which represents the "signal") is zero. On the other hand, a perfectly asymmetric learning window, like the one used throughout this study (Equation 14), yields SNR $= \frac{\Delta w_{ij}}{\mathrm{std}(\Delta w_{ij}^k)}$, because $\Delta w_{ij}^k = -\Delta w_{ji}^k$. Asymmetric learning windows thus recover the classical definition of the SNR as the ratio between the average weight change and the standard deviation of the weight change.

We note that the generalized definition above can be used to calculate the SNR for arbitrary windows, such as the learning window from Bittner et al. (2017, Figure 4C).

Assuming an asymmetric window and M uncorrelated synapses with the same mean and variance of the weight change, we can write the signal-to-noise ratio as

$$\mathrm{SNR} = \frac{M \cdot \Delta w_{ij}}{\sqrt{\mathrm{var}\left(\sum_{k=1}^{M} \Delta w_{ij}^k\right)}} = \sqrt{M} \frac{\Delta w_{ij}}{\mathrm{std}(\Delta w_{ij}^k)} \,,$$

because the variance of the sum can be decomposed into the sum of variances and covariances. All covariances are zero because synapses are uncorrelated. This leaves a sum of M variances, which are identical. Therefore, the standard deviation, and consequently also the SNR, scale with $\sqrt{M}$.

Numerical simulations

To numerically simulate the synaptic weight change, spikes were generated by inhomogeneous Poisson processes with rate functions according to Equation 11. For every spike pair, the contribution to the weight change was calculated according to Equation 14. We repeated the simulations for $N = 10^4$ trials, and the mean weight change as well as the standard deviation across trials and the SNR were estimated.

References

[1] M.J. Kahana. Associative retrieval processes in free recall. *Mem Cognit*, 24(1):103–109, 1996.

[2] N.J. Fortin, K.L. Agster, and H.B. Eichenbaum. Critical role of the hippocampus in memory for sequences of events. *Nat Neurosci*, 5(5):458–462, 2002.

[3] H. Lehn, H.-A. Steffenach, N.M. van Strien, D.J. Veltman, M.P. Witter, and A.K. Håberg. A specific role of the human hippocampus in recall of temporal sequences. *J Neurosci*, 29(11):3475–3484, 2009.

[4] P.J. Drew and L.F. Abbott. Extending the effects of spike-timing-dependent plasticity to behavioral timescales. *Proc Natl Acad Sci U S A*, 103(23):8876–8881, 2006.

[5] J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Res*, 34:171–175, 1971.

[6] H. Markram, J. Lübke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275:213–215, 1997.

[7] G. Bi and M. Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 18(24):10464–10472, 1998.

[8] J. O'Keefe and M. L. Recce. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3):317–330, 1993.

[9] G. Buzsáki. Theta oscillations in the hippocampus. *Neuron*, 33(3):325–340, 2002.

[10] R. C. Schmidt, K. Diba, C. Leibold, D. Schmitz, G. Buzsáki, and R. Kempter. Single-trial phase precession in the hippocampus. *J Neurosci*, 29(42):13232–13241, 2009.

[11] E.T. Reifenstein, R. Kempter, S. Schreiber, M.B. Stemmler, and A.V.M. Herz. Grid cells in rat entorhinal cortex encode physical space with independent firing fields and phase precession at the single-trial level. *Proc Natl Acad Sci U S A*, 109(16):6301–6306, 2012.

[12] S. Cheng and L.M. Frank. New experiences enhance coordinated neural activity in the hippocampus. *Neuron*, 57(297):303–313, 2008.

[13] W. E. Skaggs, B. L. McNaughton, M. A. Wilson, and C.A. Barnes. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–172, 1996.

[14] M.V. Tsodyks, W.E. Skaggs, T.J. Sejnowski, and B.L. McNaughton. Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. *Hippocampus*, 6:271–280, 1996.

[15] G. Bi and M. Poo. Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu Rev Neurosci*, 24:139–166, 2001.

[16] R.C. Froemke, M. Poo, and Y. Dan. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature*, 434:221–225, 2005.

[17] G.M. Wittenberg, and S. S.-H. Wang. Malleability of spike-timing-dependent plasticity at the CA3-CA1 synapse *J Neurosci*, 26(24):6610–6617, 2006.

[18] L.F. Abbott and S.B. Nelson. Synaptic plasticity: taming the beast. *Nat Neurosci*, 3:1178–1183, 2000.

[19] S. Lewandowsky and B. B. Murdock. Memory for serial order. *Psychol Rev*, 96(1):25–57, 1989.

[20] N. Sato and Y. Yamaguchi. Memory encoding by theta phase precession in the hippocampal network. *Neural Comput*, 15(10):2379–2397, 2003.

[21] C. Leibold and R. Kempter. Memory capacity for sequences in a recurrent network with biological constraints. *Neural Comput*, 18(4):904–941, 2006.

[22] E. Shen, R. Wang, and Z. Zhang. Theta phase precession enhance single trial learning in an STDP network. *Adv Cog Neurodyn ICCN*, 109–114, 2007.

[23] S. Cheng. The CRISP theory of hippocampal function in episodic memory. *Front Neural Circuits*, 7:88, 2013.

[24] N. Chenkov, H. Sprekeler, and R. Kempter. Memory replay in balanced recurrent networks. *PLoS Comput Biol*, 13(1):e1005359, 2017.

[25] P. Malerba and M. Bazhenov. Circuit mechanisms of hippocampal reactivation during sleep. *Neurobiol Learn Mem*, 160:98–107, 2019.

[26] M. Gillett, U. Pereira, and N. Brunel. Characteristics of sequential activity in networks with temporally asymmetric Hebbian learning. *Proc Natl Acad Sci U S A*, 117(47):29948–29958, 2020.

[27] K. Diba and G. Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nat Neurosci*, 10(10):1241–1242, 2007.

[28] N. Schuck and Y. Niv. Sequential replay of nonspatial task states in the human hippocampus. *Science*, 364 (6447):eaaw5181, 2019.

[29] D.O. Hebb. *The Organization of Behavior*. Wiley: New York, 1949.

[30] R. Kempter, W. Gerstner, and J.L. van Hemmen. Hebbian learning and spiking neurons. *Phys Rev E Stat Nonlin Soft Matter Phys*, 59:4498–4514, 1999.

[31] C. Geisler, K. Diba, E. Pastalkova, K. Mizuseki, S. Royer, and G. Buzsáki. Temporal delays among place cells determine the frequency of population theta oscillations in the hippocampus. *Proc Natl Acad Sci U S A*, 107(17):7957–7962, 2010.

[32] G. Dragoi and G. Buzsáki. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron*, 50(1):145–157, 2006.

[33] T. D'Albis, J. Jaramillo, H. Sprekeler, and R. Kempter. Inheritance of hippocampal place fields through Hebbian learning: Effects of theta modulation and phase precession on structure formation. *Neural Comput*, 27(8):1624–1672, 2015.

[34] J.-P. Pfistner, and W. Gerstner. Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci*, 26(38):9673–9682, 2006.

[35] K.C. Bittner, A.D. Milstein, C. Grienberger, S. Romani, and J.C. Magee. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*, 357(6355):1033–1036, 2017.

[36] T. Masquelier, E. Hugues, G. Deco, and S.J. Thorpe. Oscillations, phase-of-firing coding, and spike timing-dependent plasticity: an efficient learning scheme. *J Neurosci*, 29(43):13484–13493, 2009.

[37] J. Huxter, N. Burgess, and J. O'Keefe. Independent rate and temporal coding in hippocampal pyramidal cells. *Nature*, 425:828–832, 2003.

[38] H.B. Eichenbaum. Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci*, 15(11):732–744, 2014.

[39] M. Hasselmo. Neuromodulation: acetylcholine and memory consolidation. *Trends Cog Sci*, 3(9):351–359, 1999.

[40] J.C. Zhang, P.M. Lau, and G.Q. Bi. Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. *Proc Natl Acad Sci U S A*, 106(31):13028–13033, 2009.

[41] N. Ulanovsky and C.F. Moss. Hippocampal cellular and network activity in freely moving echolocating bats. *Nat Neurosci*, 10(2):224–233, 2007.

[42] N.A. Herweg and M.J. Kahana. Spatial representations in the human brain. *Front Hum Neurosci*, 12(297), 2018.

[43] C. Leibold, A. Gundlfinger, R. Schmidt, K. Thurley, D. Schmitz, and R. Kempter. Temporal compression mediated by short-term synaptic plasticity. *Proc Natl Acad Sci U S A*, 105:4417–4422, 2008.

[44] J. Haam, J. Zhou, G. Cui, and J.L. Yakela. Septal cholinergic neurons gate hippocampal output to entorhinal cortex via oriens lacunosum moleculare interneurons. *Proc Natl Acad Sci U S A*, 115(8):E1886–E1895, 2018.

[45] N. Rebola, M. Carta, and C. Mulle. Operation and plasticity of hippocampal CA3 circuits: implications for memory encoding. *Nat Rev Neurosci*, 18(4):208–220, 2017.

[46] A. Chadwick, M.Cw. van Rossum, M.F. Nolan. Flexible theta sequence compression mediated via phase precessing interneurons. *Elife*, 5:e20349, 2016.

[47] F.S. Chance. Hippocampal phase precession from dual input components. *J Neurosci*, 32(47):16693–16703, 2012.

[48] J. Jaramillo, R. Schmidt, and R. Kempter. Modeling inheritance of phase precession in the hippocampal formation. *J Neurosci*, 34(22):7715–7731, 2014.

[49] J.C. Magee. Dendritic mechanisms of phase precession in hippocampal CA1 pyramidal neurons. *J Neurophysiol*, 86(1):528–532, 2001.

[50] K.D. Harris, D.A. Henze, H. Hirase, X. Leinekugel, G. Dragoi, A. Czurkó, and G. Buzsáki. Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature*, 417(6890):738–741, 2002.

[51] M.R. Mehta, A.K. Lee, and M.A. Wilson. Role of experience and oscillations in transforming a rate code into a temporal code. *Nature*, 417(6890):741–746, 2002.

[52] K. Thurley, C. Leibold, A. Gundlfinger, D. Schmitz, and R. Kempter. Phase precession through synaptic facilitation. *Neural Comput*, 20:1285–1324, 2007.

[53] J. Jaramillo and R. Kempter. Phase precession: a neural code underlying episodic memory? *Curr Opin Neurobiol*, 43:130–138, 2017.

[54] E.V. Lubenov, A.G. Siapas. Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246): 534–539, 2009.

[55] J. Patel, S. Fujisawa, A. Berényi, S. Royer, G. Buzsáki. Traveling theta waves along the entire septotemporal axis of the hippocampus. *Neuron*, 75(3):410–417, 2012.

[56] C. Leibold and M.M. Monsalve-Mercado. Traveling theta waves and the hippocampal phase code. *Sci Rep*, 7(1):7678, 2017.

[57] L. Muller, F. Chavane, J. Reynolds, and T.J. Sejnowski. Cortical travelling waves: mechanisms and computational principles. *Nat Rev Neurosci*, 19(5):255–268, 2018.

[58] P.J. Sjöström, G.G. Turrigiano, and S.B. Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32:1149–1164, 2001.

[59] M.A. Wilson and B.L. McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994.

[60] Z. Nadásdy, H. Hirase, A. Czurko, J. Csicsvari, and G. Buzsáki. Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci*, 19(21):9497–9507, 1999.

[61] A. Peyrache, M. Khamassi, K. Benchenane, S.I. Wiener, and F.P. Battaglia. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat Neurosci*, 12(7):919–926, 2009.

[62] T.J. Davidson, F. Kloosterman, and M.A. Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, 2009.

[63] M.M. Gauy, J. Lengler, H. Einarsson, F. Meier, F. Weissenberger, M.F. Yanik, and A. Steger. A hippocampal model for behavioral time acquisition and fast bidirectional replay of spatio-temporal memory sequences. *Front Neurosci*, 12:961, 2020.

[64] E.R. Wood, P.A. Dudchenko, R.J. Robitsek, H.B. Eichenbaum, D. Kumaran, and E.A. Maguire. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, 27(3):623–633, 2000.

[65] N.M. Long and M.J. Kahana. Hippocampal contributions to serial-order memory. *Hippocampus*, 29(3): 252–259, 2019.

Acknowledgements

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation; Grants GRK 1589/2, SPP 1665, SFB 1315 – project-ID 327654276) and the German Federal Ministry for Education and Research (BMBF; Grant 01GQ1705). We thank Ikhwan Khalid, Lukas Kunz, Natalie Schieferstein, Tiziano D'Albis, Paul Pfeiffer, and Adam Wilkins for helpful discussions and feedback on the manuscript. ETR and RK

designed and performed the research, wrote and discussed the manuscript. ETR and RK declare no conflict of interest.