

Coding and learning of behavioral sequences

Ofer Melamed^{1,4}, Wulfram Gerstner^{1,2}, Wolfgang Maass³, Misha Tsodyks⁴ and Henry Markram¹

¹Brain Mind Institute, EPFL, Lausanne 1015, Switzerland

²Laboratory of Computational Neuroscience, EPFL 1015, Lausanne, Switzerland

³Institute for Theoretical Computer Science, Technische Universität Graz, Austria

⁴Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel

A major challenge to understanding behavior is how the nervous system allows the learning of behavioral sequences that can occur over arbitrary timescales, ranging from milliseconds up to seconds, using a fixed millisecond learning rule. This article describes some potential solutions, and then focuses on a study by Mehta *et al.* that could contribute towards solving this puzzle. They have discovered that an experience-dependent asymmetric shape of hippocampal receptive fields combined with oscillatory inhibition can serve to map behavioral sequences on a fixed timescale.

Behavioral sequences of events span several orders of magnitude of timescales. A piano player is able to play rapid scales or repetitions of a single tone at a rate of 12 or more keys per second, probably using an internal sequence of motor commands that is even faster. However, when playing a slow tune, commands are executed ten or twenty times slower. Such different temporal sequences should ideally all be learned using a synaptic learning rule that operates on a fixed neuronal timescale of milliseconds [1–3]. How then, is learning of behavioral sequences across arbitrary timescales possible? Several studies have recently addressed this issue from different points of view, and have suggested potential solutions.

Coding sequences at millisecond resolution

One potential solution of the problem could be that all sequences, fast or slow, are in fact stored and retrieved at a millisecond resolution. Synfire chains (SFC) [4] – dominantly feedforward multi-layered architecture (chains) in which spiking activity can propagate in a synchronous wave of neuronal firing (a pulse packet) from one layer of the chain to the next – provide a conceptual framework for such an approach. The chain of neuronal activity evolves on a millisecond scale and, hence, on the natural timescale of spike-time-dependent plasticity. Because all information is represented by a sequence of firing neuronal times, the SFC would automatically solve the problem of sequence learning using a fixed-timescale learning rule. The fastest behavioral sequence would be one in which each time step of the neuronal chain corresponds to one behavioral event. A slow behavioral sequence would then simply be represented by several neuronal time steps per behavioral event. The problem with such an approach is

that the representation of slow sequences of events (e.g. five events in one second) is rather expensive, because it requires hundreds of transitions that need to be precisely stored in neuronal connections [4–6].

A recent study suggested that a mechanism closely related to neuronal SFCs is implemented in the premotor area of song-birds [7]. Neurons projecting from the nucleus hyperstriatum ventralis pars caudalis (HVC) to the forebrain robust nucleus of the archistriatum (RA) are active in a tight temporal sequence or ‘clock’. In contrast to the SFC concept, however, each neuron emits not a single spike but a burst of three to four spikes within ~10 ms. Learning a song could then occur in the synapses connecting HVC to RA. It is worth noting, however, that in contrast to piano players who can play the same tune at a different ‘tempo’, birds usually generate their songs at a fixed speed.

Coding slow sequences

For slow sequences, it would be preferable to encode only transitions between relevant behavioral events that typically occur on the timescale of hundreds of milliseconds or slower. One theoretical possibility to achieve this would be to work with a sequence of memory items. In this case, each item corresponds to one behavioral event and is represented by the firing rate of a set of active neurons. The transitions from one item to the next could then be triggered by adaptation of the active neurons [8], by synaptic depression of synapses between active neurons [9] or by proprioceptive feedback during behavior [10]. Learning is then possible by any form of asymmetric Hebbian plasticity [10,11].

Another theoretical possibility would be to encode behavioral sequences in neurons that are part of a slow loop with an oscillation frequency of, say, 10 Hz. In mammals, learning and retrieval of behavioral sequences probably involves the cerebellum. Inhibitory rebound in deep cerebellar nuclei can slow down the neuronal timescale to ~100 ms. A loop connecting deep cerebellar nuclei, the cerebellar cortex and the inferior olive could then provide a clock with ~100 ms cycle time which would make it easy to connect to behavioral timescales [12].

Sequences of hippocampal place cells

Phase codes in the hippocampus offer yet another interesting mechanism for sequence encoding. Recently,

Corresponding author: Henry Markram (henry.markram@epfl.ch).

several lines of evidence in the hippocampus have pointed towards an interesting and very simple cellular mechanism that could allow learning of sequences of locations with some invariance in the speed of movement using a fixed learning rule on the millisecond timescale. The first crucial discovery was when O'Keefe and colleagues found neurons responding only at specific locations of a running track – place cells [13]. O'Keefe and colleagues then realized that while the rat moves through the preferred location, the firing times preferably occur at a well defined phase with respect to a background oscillation at theta-frequencies [14]. Furthermore, O'Keefe and colleagues found that the phase of firing decreases as the rat moves deeper into the place field. The phase precession therefore means that the phase at which a neuron fires with respect to the theta cycle carries information about where the rat is [14]. Skaggs and colleagues pointed out that this effect means that neurons with overlapping place fields, following each other along the path of the rat, will fire in a specific temporal order during each theta cycle – in other words, this represents a fast replay of sequential firing in space within one theta cycle [15] (Figure 1c). Several researchers have since speculated that this could be involved in prediction, planning, sequence learning and other functions [16–21].

Insights into the mechanism of phase precession

The mechanisms for this phase precession have been debated extensively. A breakthrough came with a study by Mehta and colleagues [22–24] when they studied the shape of the place field. They found that after several trials, cells slowly increased the firing rate while the rat moved through the place field, to reach a maximum just before the end of the field. The results can be summarized as a highly asymmetrical place field. Mehta and colleagues then realized that asymmetric excitation, which is suggested by the asymmetrical place field, combined with an oscillatory inhibition, could account for the phase precession (Figure 1a). An oscillatory inhibition therefore transforms a rate code into a temporal code where the firing of activity in different place cells can be ordered on different phases of a single theta cycle. In addition, this mechanistic explanation predicts that with increasing excitation the precision of phase coding decreases [24]. Although the onset of firing of different place cells is nicely ordered along the phase axis, the most activated place cells fire longer spike bursts than those that are less excited. This could have a strong influence on synaptic plasticity [25,26]. Spike train adaptation might, however, be important to limit the burst duration, despite

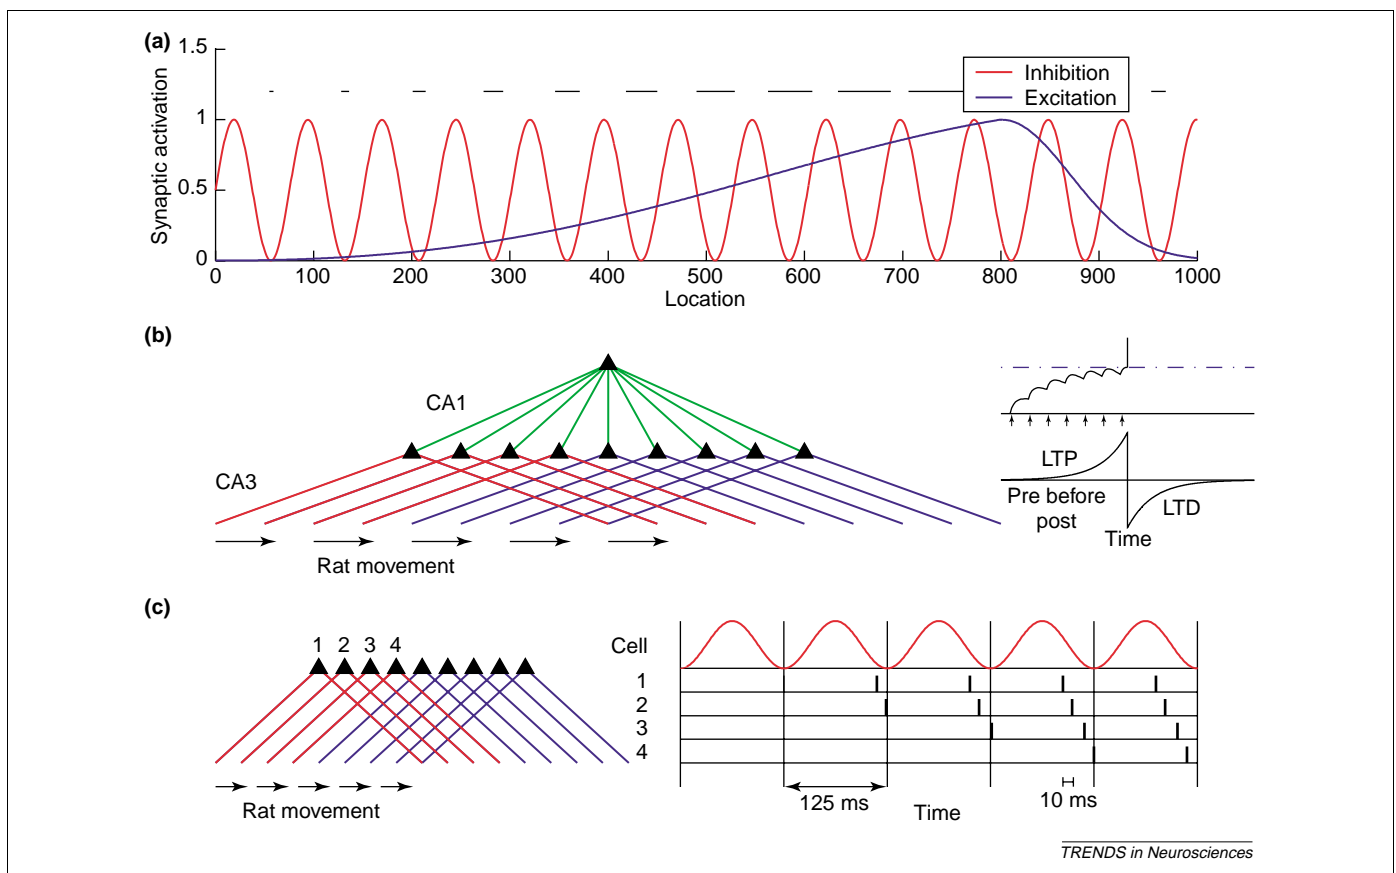


Figure 1. Timing and learning. **(a)** Conversion of rate code to temporal code. Phase precession generated by superimposing oscillatory inhibition (red) with an asymmetrical receptive field (blue). The neuron can fire only if excitation exceeds inhibition, implying that the phase at which firing starts decreases and the maximal duration of spike bursts (indicated by horizontal bars) increases. **(b)** Asymmetric receptive fields caused by asymmetric Hebbian learning. Because a postsynaptic neuron in the CA1 hippocampus fires only after several excitatory postsynaptic potentials (EPSPs; arrows in the graph) have accumulated, connections from the CA3 place cells that are stimulated first are potentiated [i.e. they undergo long-term potentiation (LTP)]. Green lines represent inputs from CA3 cells to a single CA1 cell; red and blue lines represent boundaries of the place field for each CA3 place cell. The lower graph shows the window of spike-timing-dependent plasticity (STDP) with LTP induced at synapses whose presynaptic spikes arrive before the postsynaptic firing time ('Pre before post'). See also Ref. [30]. **(c)** Compression of firing sequence. Because of phase precession, the activation sequence of place cells 1 to 4 as a rat traverses the corresponding place fields during ~0.5 s movement is repeated within a single theta cycle (red trace). As a result, the temporal sequence is played on a timescale of 10 ms, which is suitable for STDP.

increased excitation. As a result, spike-train adaptation allows peak firing rates before the peak of an oscillatory excitation [27,28].

Asymmetric place fields and asymmetric learning

The asymmetry of place fields could be the result of asymmetric Hebbian plasticity that strengthens synapses at which the presynaptic neuron fires before the postsynaptic one, and weakens those at which the spike timing is reversed [1–3]. For the development of asymmetric place fields (Figure 1b), rate coding is sufficient [10,11,29,30] and a precise temporal code is not required. The development of asymmetric place fields [24,30] or, similarly, of visual receptive fields [31,32] could therefore be interpreted as a system-level signature of asymmetric Hebbian plasticity.

In the context of sequence learning, a second aspect is equally important. Phase precession of hippocampal place cells leads to a temporal code in which slow sequences are temporally compressed within one theta cycle (Figure 1c). Because the spikes of different neurons are ordered within one cycle on a millisecond timescale, phase coding provides the precise timing required in spike-timing-dependent plasticity (STDP) to determine the magnitude and direction of synaptic changes. Given phase coding, STDP would then allow learning of transitions between places (i.e. provide a substrate for behavioral sequences).

Open questions

There are however, several problems that need to be addressed. Mehta and colleagues found the shift from a symmetrical to asymmetrical place field during training, but they also found that this development of asymmetry was completely reset the following day [30]. Combining a symmetrical receptive field with an oscillatory inhibition would lead to a phase precession and then a recession, the result being that the spiking at a particular phase could not reliably represent the location. The problem is that, although Mehta and colleagues report a resetting of the symmetry, phase precession is always seen in previous reports [33–35]. Does that mean that phase precession is generated by yet a different learning mechanism? Even if this were the case, it raises yet another problem: is the sequence learned before or after the asymmetry of the place field evolves? The neural trick of temporal compression seems to be useful for a learning rule on the millisecond scale only if receptive fields are already asymmetric.

A completely different line of research suggests that the puzzle of bridging timescales is potentially less problematic than previously thought because, apart from the millisecond timescale of neural firing, neural dynamics is also characterized by other time constants, ranging from milliseconds up to seconds. Some proposals [36,37] suggest analyzing the problem of temporal information encoding from the point of view of readout neurons (e.g. in CA1) that have learnt to extract at any moment in time information from the current state of a neural microcircuit (e.g. in CA3). Because of the long cellular, synaptic and network time constants of neural circuits [38,39], the diversity of neural components and the asymmetric sparse connection patterns, information about a rapid and high-dimensional

input circulates in the microcircuit and is available to readout neurons at a much later time. Thus, bridging timescales could be natural for neural microcircuits and readout neurons can learn to recognize relevant information using simple perception-like learning rules [37].

In summary, an interesting and very simple neural trick has been revealed, in which asymmetrical excitation combined with oscillatory inhibition can order the precise timing of spikes in neurons with different optimal receptive fields in space and in time. This squeezing of activity across neurons into a standardized time frame could allow a relatively fixed learning rule to scale synapses according to their relative spatial location in a brain region and the temporal sequence of their activation. This mechanism might therefore be general for receptive fields in sensory cortices whenever periodic background signals are available.

References

- 1 Markram, H. *et al.* (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215
- 2 Zhang, L.I. *et al.* (1998) A critical window for cooperation and competition among developing retinotectal synapses. *Nature* 395, 37–44
- 3 Bi, G. and Poo, M. (2001) Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu. Rev. Neurosci.* 24, 139–166
- 4 Abeles, M. (1991) *Corticonics: Neural Circuits of the Cerebral Cortex*, Cambridge University Press
- 5 Aviel, Y. *et al.* (2002) Synfire chain in a balanced network. *Neurocomputing* 44, 285–292
- 6 Gerstner, W. *et al.* (1993) Why spikes? Hebbian learning and retrieval of time-resolved excitation patterns. *Biol. Cybern.* 69, 503–515
- 7 Hahnloser, R.H. *et al.* (2002) An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419, 65–70
- 8 van Vreeswijk, C. and Hansel, D. (2001) Patterns of synchrony in neural networks with spike adaptation. *Neural Comput.* 13, 959–992
- 9 Markram, H. *et al.* (1998) Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Natl. Acad. Sci. U. S. A.* 95, 5323–5328
- 10 Abbott, L.F. and Blum, K.I. (1996) Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex* 6, 406–416
- 11 Roberts, P.D. (1999) Computational consequences of temporally asymmetric learning rules: I. Differential Hebbian learning. *J. Comput. Neurosci.* 7, 235–246
- 12 Kistler, W.M. *et al.* (2000) Time window control: a model for cerebellar function based on synchronization, reverberation, and time slicing. *Prog. Brain Res.* 124, 275–297
- 13 O'Keefe, J. and Dostrovsky, J. (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175
- 14 O'Keefe, J. and Recce, M.L. (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317–330
- 15 Skaggs, W.E. *et al.* (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6, 149–172
- 16 Hopfield, J.J. (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376, 33–36
- 17 Tsodyks, M.V. *et al.* (1996) Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. *Hippocampus* 6, 271–280
- 18 Wallenstein, G.V. and Hasselmo, M.E. (1997) GABAergic modulation of hippocampal population activity: sequence learning, place field development, and the phase precession effect. *J. Neurophysiol.* 78, 393–408
- 19 Bose, A. *et al.* (2000) A temporal mechanism for generating the phase precession of hippocampal place cells. *J. Comput. Neurosci.* 9, 5–30
- 20 Jensen, O. and Lisman, J.E. (1996) Hippocampal CA3 region predicts

- memory sequences: accounting for the phase precession of place cells. *Learn. Mem.* 3, 279–287
- 21 Jensen, O. and Lisman, J.E. (1996) Theta/gamma networks with slow NMDA channels learn sequences and encode episodic memory: role of NMDA channels in recall. *Learn. Mem.* 3, 264–278
- 22 Mehta, M.R. *et al.* (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. U. S. A.* 94, 8918–8921
- 23 Mehta, M.R. (2001) Neuronal dynamics of predictive coding. *Neuroscientist* 7, 490–495
- 24 Mehta, M.R. *et al.* (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417, 741–746
- 25 Taschenberger, H. *et al.* (2002) Optimizing synaptic architecture and efficiency for high-frequency transmission. *Neuron* 36, 1127–1143
- 26 Buzsaki, G. *et al.* (2002) Homeostatic maintenance of neuronal excitability by burst discharges *in vivo*. *Cereb. Cortex* 12, 893–899
- 27 Fuhrmann, G. *et al.* (2002) Spike frequency adaptation and neocortical rhythms. *J. Neurophysiol.* 88, 761–770
- 28 Harris, K.D. *et al.* (2002) Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature* 417, 738–741
- 29 Gerstner, W. and Abbott, L.F. (1997) Learning navigational maps through potentiation and modulation of hippocampal place cells. *J. Comput. Neurosci.* 4, 79–94
- 30 Mehta, M.R. *et al.* (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25, 707–715
- 31 Fu, Y.X. *et al.* (2002) Temporal specificity in the cortical plasticity of visual space representation. *Science* 296, 1999–2003
- 32 Engert, F. *et al.* (2002) Moving visual stimuli rapidly induce direction sensitivity of developing tectal neurons. *Nature* 419, 470–475
- 33 Czurko, A. *et al.* (1999) Sustained activation of hippocampal pyramidal cells by ‘space clamping’ in a running wheel. *Eur. J. Neurosci.* 11, 344–352
- 34 Ekstrom, A.D. *et al.* (2001) NMDA receptor antagonism blocks experience-dependent expansion of hippocampal ‘place fields’. *Neuron* 31, 631–638
- 35 Magee, J.C. (2001) Dendritic mechanisms of phase precession in hippocampal CA1 pyramidal neurons. *J. Neurophysiol.* 86, 528–532
- 36 Buonomano, D.V. and Merzenich, M. (1999) A neural network model of temporal code generation and position-invariant pattern recognition. *Neural Comput.* 11, 103–116
- 37 Maass, W. *et al.* (2002) Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* 14, 2531–2560
- 38 Gupta, A. *et al.* (2000) Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science* 287, 273–278
- 39 Thomson, A.M. and Deuchars, J. (1994) Temporal and spatial properties of local circuits in neocortex. *Trends Neurosci.* 17, 119–126

0166-2236/\$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.tins.2003.10.014

Research Focus Response

Response to Melamed *et al.*: Coding and learning of behavioral sequences – open questions and potential solutions

Mayank R. Mehta¹, Albert K. Lee² and Matthew A. Wilson²

¹Department of Neuroscience, Brown University, Providence, RI 02912, USA

²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

What are the neuronal mechanisms involved in learning behavioral sequences? This is a difficult and important question. Several recent studies have uncovered unforeseen problems involved in solving this question, and have offered novel solutions. In our recent work we focused on one of the problems involved – namely, the importance of replaying behavioral sequences at the relatively fast time scales of synaptic plasticity [1]. Neurons that encode events within a behavioral sequence, such as walking from place A to place B, are typically active over periods of a second or longer. However, to learn the behavioral sequence rapidly via NMDA-dependent synaptic plasticity these neurons should be activated in the same temporal order within ~ 10 ms. We presented a computational mechanism that could accomplish this task, made several predictions using this mechanism, and verified these predictions experimentally using data from hippocampal place cells [1]. In an accompanying article, Melamed *et al.* have raised several interesting open questions raised by our work [2]. Here, we point to potential solutions for three of these questions.

(i) ‘Combining a symmetrical receptive field with an oscillatory inhibition would lead to a phase precession and

then a recession, the result being that the spiking at a particular phase could not reliably represent the location’.

This would indeed be a problem if the position were encoded by the phase of a single symmetric place field. However, this is clearly not the case because a large number of hippocampal neurons are active at any given location. The position can be unambiguously encoded by the phases of symmetric receptive fields across a population of neurons with overlapping place fields. Indeed, several cortical receptive fields have a bell-shaped or symmetrical firing-rate profile as a function of certain physical variable. Whereas the firing rate of an individual cell provides ambiguous information about the physical variable, a population code consisting of firing rates of a large number of neurons with overlapping symmetric receptive fields is known to provide unambiguous information. Ditto for the temporal code. Of course, asymmetric firing rate and asymmetric temporal codes would remove the ambiguity in encoding a physical variable at a single-cell level using a rate or a temporal code, respectively.

(ii) ‘The problem is that, although Mehta and colleagues report a resetting of the symmetry, phase precession is always seen in previous reports’.

Corresponding author: Mayank R. Mehta (mayank_mehta@brown.edu).