



Recall Performance Improvement in a Bio-Inspired Model of the Mammalian Hippocampus

Nikolaos Andreakos¹ , Shigang Yue¹ ,
and Vassilis Cutsuridis^{1,2} 

¹ School of Computer Science, University of Lincoln, Lincoln, UK
{nandreakos, syue, vcutsuridis}@lincoln.ac.uk

² Lincoln Sleep Research Center, University of Lincoln, Lincoln, UK

Abstract. Mammalian hippocampus is involved in short-term formation of declarative memories. We employed a bio-inspired neural model of hippocampal CA1 region consisting of a zoo of excitatory and inhibitory cells. Cells' firing was timed to a theta oscillation paced by two distinct neuronal populations exhibiting highly regular bursting activity, one tightly coupled to the trough and the other to the peak of theta. To systematically evaluate the model's recall performance against number of stored patterns, overlaps and 'active cells per pattern', its cells were driven by a non-specific excitatory input to their dendrites. This excitatory input to model excitatory cells provided context and timing information for retrieval of previously stored memory patterns. Inhibition to excitatory cells' dendrites acted as a non-specific global threshold machine that removed spurious activity during recall. Out of the three models tested, 'model 1' recall quality was excellent across all conditions. 'Model 2' recall was the worst. The number of 'active cells per pattern' had a massive effect on network recall quality regardless of how many patterns were stored in it. As 'active cells per pattern' decreased, network's memory capacity increased, interference effects between stored patterns decreased, and recall quality improved. Key finding was that increased firing rate of an inhibitory cell inhibiting a network of excitatory cells has a better success at removing spurious activity at the network level and improving recall quality than increasing the synaptic strength of the same inhibitory cell inhibiting the same network of excitatory cells, while keeping its firing rate fixed.

Keywords: Associative memories · Neural information processing · Brain · Inhibition

1 Introduction

Memory is our most precious faculty. The case of Henry Molaison (the infamous 'HM' patient) has taught us a lot about what happens when we cannot store memories. Without memory we are unable to remember our past experiences and our loved ones, while still being able to think about the future. Without memory we cannot learn

anything new. Associative memory is the ability to learn and remember the relationship between items, places, events, and/or objects which may be unrelated [1].

Hippocampus, the site of short-term storage of declarative memories [2], is one of the most studied brain areas yielding a wealth of knowledge of cell types and their anatomical, physiological, synaptic, and network properties [3]. Cells in various hippocampal regions have been hypothesized to compute information differently. Regions CA3 and CA1 have also been implicated in auto- and hetero-association (storage) of declarative memories, respectively [4].

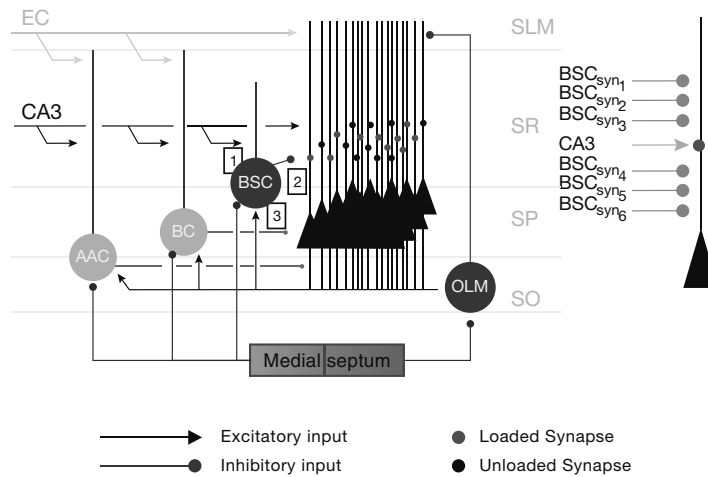


Fig. 1. Associative neural network model of region CA1 of the hippocampus and CA1-PC model with one excitatory (CA3) and six inhibitory (BSC) synaptic contacts on its SR dendrites. During retrieval only PC, BSC, and OLM cells are active. AAC and BC are inactive due to strong medial septum inhibition. BSC and PC are driven on their SR dendrites by a strong CA3 excitatory input, which presented the contextual information. Red circles on PC dendrites represent loaded synapses, whereas black circles on PC dendrites represent unloaded synapses. EC: Entorhinal cortical input; CA3: Schaffer collateral input; AAC: Axo-axonic cell; BC: basket cell; BSC: bistratified cell; OLM: oriens lacunosum-moleculare cell; SLM: stratum lacunosum moleculare; SR: stratum radiatum; SP: stratum pyramidale; SO: stratum oriens.

In 2010 a bio-inspired microcircuit model of region CA1 was introduced that controlled for itself the storage and recall of patterns of information arriving at high rates [5]. The model was based upon the biological details were then known about the hippocampal neural circuit [6, 7]. The model explored the functional roles of somatic, axonic and dendritic inhibition in the encoding and retrieval of memories. It showed how theta modulated inhibition separated encoding and retrieval of memories into two functionally independent processes. It showed how somatic inhibition allowed generation of dendritic calcium spikes that promoted synaptic long-term plasticity (LTP), while minimizing cell output. Proximal dendritic inhibition controlled both cell output and suppressed dendritic calcium spikes, thus preventing LTP, whereas distal dendritic inhibition removed interference from spurious memories during recall. The mean recall quality of the model was tested as function of memory patterns stored. Recall dropped

as more patterns were encoded due to interference between previously stored memories.

Here, we systematically investigate the biophysical mechanisms of this bio-inspired neural network model of region CA1 of the hippocampus [5] to improve its memory capacity and recall performance. In particular, we examine how selective modulation of feedforward/feedback excitatory/inhibitory pathways targeting inhibitory and excitatory cells may influence the thresholding ability of dendritic inhibition to remove at the network level spurious activities, which may otherwise impair the recall performance of the network, and improve its mean recall quality as more and more overlapping memories are stored.

2 Materials and Methods

2.1 Neural Network Model

Figure 1 depicts the simulated neural network model of region CA1 of the hippocampus. The model consisted of 100 excitatory cells (pyramidal cells (PC)) and four types of inhibitory cells: 1 axo-axonic cell (AAC), 2 basket cells (BC), 1 bistratified (BSC) and 1 oriens lacunosum-moleculare (OLM) cell. Simplified morphologies including the soma, apical and basal dendrites and a portion of the axon were used for each cell type. The biophysical properties of each cell were adapted from cell types reported in the literature, which were extensively validated against experimental data in [8–12]. Using known physical properties and effects of cell structures is a more efficient way to examine scientific hypothesis compare to blind computational optimization. The core of our research was biological properties and mechanisms because by we obtained a better understanding on how these mechanisms affected the whole circuit and gained some insightful intuitions. The complete mathematical formalism of the model has been described elsewhere [5]. Schematic representations of model cells can be found in [13]. The dimensions of the somatic, axonic and dendritic compartments of model cells, the parameters of all passive and active ionic conductances, synaptic waveforms and synaptic conductances can be found in [13]. All simulations were performed using NEURON [14] running on a PC with four CPUs under Windows 10.

2.2 Inputs

Network was driven by an excitatory CA3 input and an inhibitory medial septum (MS) input. The excitatory input was modelled as the firing of 20 out of 100 CA3 pyramidal cells at an average gamma frequency of 40 Hz (spike trains only modelled and not the explicit cells). PCs, BCs, AACs, BSCs in our network received excitatory input in their proximal-to-soma dendrites. The inhibitory input was modelled with the rhythmic firing of two opponent processing populations of 10 inhibitory cells each firing at opposite phases of a theta cycle (180° out of phase) [15]. Each such cell output was modelled as bursts of action potentials using a presynaptic spike generator. Each spike train consisted of bursts of action potentials at a mean frequency of 8 Hz for a half-theta cycle (125 ms) followed by a half-theta cycle of silence (125 ms).




Fig. 2. Set of five memory patterns with 40% overlap between them.

Due to 8% noise in the inter-spike intervals, the 10 spike trains in each population were asynchronous. One inhibitory population input provided inhibition to BSCs and OLMs during the encoding cycle, whereas the other inhibitory population input provided inhibition to AACs and BCs during the retrieval cycle.

2.3 Network Training and Testing

The goal of this work is to test the recall performance of the model when the network had already stored memory patterns without examining the exact details of the learning process. To test the recall performance of the model the methodology described in [5] was adopted. Briefly, a memory pattern was stored by generating weight matrices based on a clipped Hebbian learning rule; these weight matrices were used to pre-specify the CA3 to CA1 PC connection weights. Without loss of generality, the input (CA3) and output (CA1) patterns were assumed to be the same, with each pattern consisting of N ($N = 5$ or 10 or 20) randomly chosen PCs (active cells per pattern) out of the population of 100. The 100 by 100 dimensional weight matrix was created by setting matrix entry (i, j) , $w_{ij} = 1$ if input PC i and output PC j are both active in the same pattern pair; otherwise weights are 0. Any number of pattern pairs could be stored to create this binary weight matrix. The matrix was applied to our network model by connecting a CA3 input to a CA1 PC with a high AMPA conductance ($g_{\text{AMPA}} = 1.5$ nS) if their connection weight was 1, or with a low conductance ($g_{\text{AMPA}} = 0.5$ nS) if their connection was 0. This approach is in line with experimental evidence that such synapses are 2-state in nature [16].

2.4 Memory Patterns

We created sets of memory patterns at different sizes (1, 5, 10, 20), percent overlaps (0%, 10%, 20%, 40%) and number of active cells per pattern (5, 10, 20). For example, a 0% overlap between N patterns in a set meant no overlap between patterns 1 and 2, 1 and 3, 1 and 4, 1 and 5, 2 and 3, 2 and 4, 2 and 5, 3 and 4, 3 and 5, and 4 and 5. A 40% overlap between 5 patterns in a set meant that $0.4 \cdot N$ cells were shared between patterns 1 and 2, a different $0.4 \cdot N$ cells were shared between patterns 2 and 3, a different $0.4 \cdot N$ cells between patterns 3 and 4, a different $0.4 \cdot N$ cells between patterns 4 and 5 and a different $0.4 \cdot N$ cells between patterns 5 and 1 (see Fig. 2). For 20 active cells per pattern that meant that a maximum of 5 patterns could be stored by a network of 100 PCs. For 10 active cells per pattern, then a maximum of 10 patterns could be stored and for 5 active cells per pattern, a maximum of 20 patterns could be stored. Similar maximum number of patterns could be stored for 10%, 20% and 40% overlap and 5, 10 and 20 active cells per pattern, respectively. In the case of 10% overlap, 5 active cells per pattern, the maximum number of stored patterns was not an integer, so this case was excluded from our simulations.

2.5 Recall Performance Measure

The recall performance metric used for measuring the distance between the recalled output pattern, B , from the required output pattern, B^* , was the correlation (i.e., degree of overlap) metric, calculated as the normalized dot product:

$$C = \frac{B \times B^*}{\left(\sum_{i=1}^{N_B} B_i \times \sum_{j=1}^{N_B} B_j^* \right)^{1/2}} \quad (1)$$

where N_B is the number of output units. The correlation takes a value between 0 (no correlation) and 1 (the vectors are identical). The higher the correlation, the better the recall performance.

2.6 Mean Recall Quality

Mean recall quality of our network model was defined as the mean value of all recall qualities estimated from each pattern presentation when an M number of patterns were already stored in the network. For example, when five patterns were initially stored in the network and pattern 1 was presented to the network during recall, then a recall quality value for pattern 1 was calculated. Repeating this process for each of the other patterns (pattern 2, pattern 3, pattern 4, and pattern 5) a recall quality value was calculated. The mean recall quality of the network was then the mean value of these individual recall qualities.

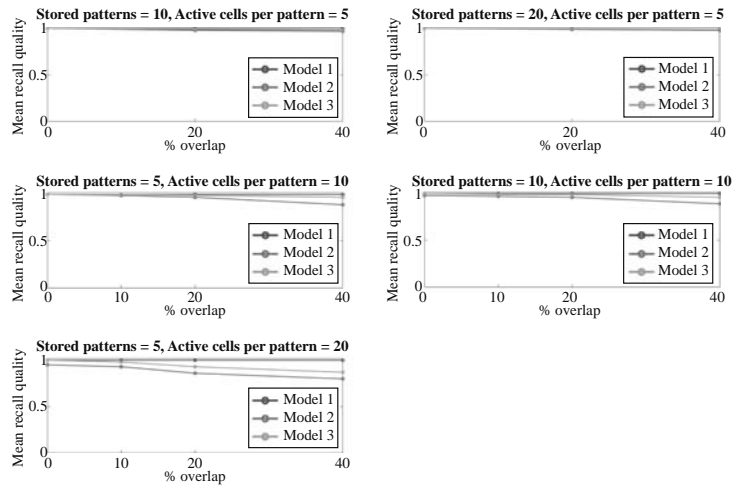


Fig. 3. Mean recall quality of ‘model 1’, ‘model 2’, and ‘model 3’ as a function of percent overlap (0%, 10%, 20%, 40%).

2.7 Model Selection

In [5], BSC inhibition to PC dendrites acted as a global non-specific threshold machine capable of removing spurious activity at the network level during recall. In [5] BSC inhibition was held constant as the number of stored patterns to PC dendrites increased. The recall quality of the model in [5] decreased as more and more memories were loaded onto the network (see Fig. 14 in [5]). To improve the recall performance of [5] we artificially modulated the synaptic strength of selective excitatory and inhibitory pathways to BSC and PC dendrites as more and more patterns were stored in the network (see Figs. 1 and 6):

1. Model 1: Increased CA3 feedforward excitation (weight) to BSC (Fig. 6A) increased the frequency of its firing rate. As a result, more IPSPs were generated in the PC dendrites producing a very strong inhibitory environment which eliminated all spurious activity.
2. Model 2: Increased BSC feedforward inhibition (weight) to PC dendrites (Fig. 6B) produced fewer IPSPs, but with greater amplitude, in the PC dendrites.
3. Model 3: Increased PC feedback excitation (weight) to BSC (Fig. 6C) had a similar effect as Model 1, but with less potency.

Comparative analysis of the above three models’ recall performance is depicted in Figs. 3, 4 and 5.

3 Results and Discussion

A set of patterns (1, 5, 10, 20) at various percent overlaps (0%, 10%, 20%, 40%) were stored by different number of ‘active cells per pattern’ (5, 10, 20) without recourse to a learning rule by generating a weight matrix based on a clipped Hebbian learning rule, and using the weight matrix to prespecify the CA3 to CA1 PC connection weights. To test recall of a previously stored memory pattern in the model, the entire associated

input pattern was applied as a cue in the form of spiking of active CA3 inputs (those belonging to the pattern) distributed within a gamma frequency time window. The cue pattern was repeated at gamma frequency (40 Hz). During the retrieval only the BSCs and OLM cells were switched on, whereas the AACs and BCs were switched off. The CA3 spiking drove the CA1 PCs plus the BSCs. The EC input, which excited the apical dendrites of PCs, AACs and BCs, was disconnected during the retrieval.

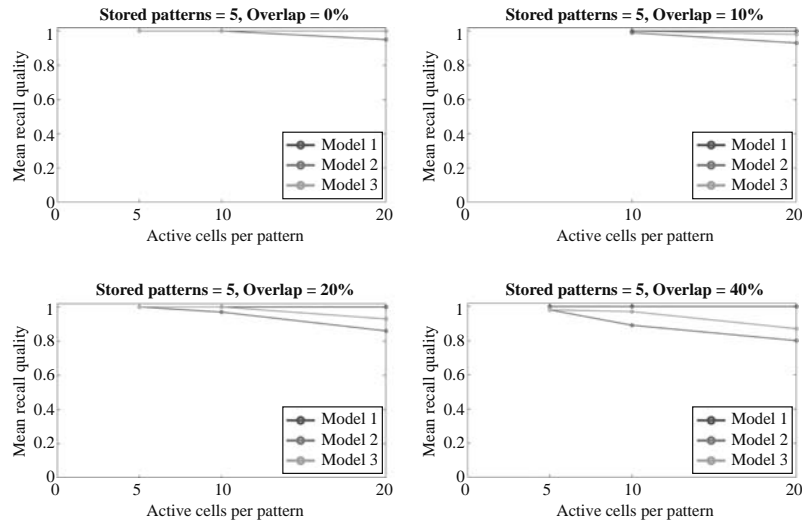


Fig. 4. Mean recall quality of ‘model 1’, ‘model 2’, and ‘model 3’ as a function of ‘active cells per pattern’. Five patterns were stored in a network of 100 PCs at 0%, 10%, 20% and 40% overlap.

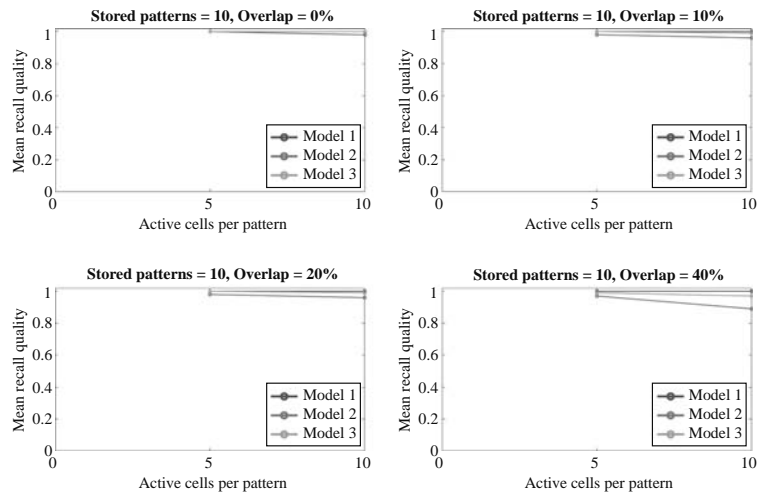


Fig. 5. Mean recall quality of ‘model 1’, ‘model 2’, and ‘model 3’ as a function of ‘active cells per pattern’. Ten patterns were stored in a network of 100 PCs at 0%, 10%, 20% and 40% overlap.

It is evident from Fig. 3 that recall performance is best for all three models (‘model 1’, ‘model 2’, ‘model 3’) when there is no overlap between patterns or when the overlap is

small (up to 10%) regardless of the number ‘active cells per pattern’ (i.e. the number of cells needed to represent a memory pattern) and patterns stored in the network. At overlaps larger than 10%, the recall performance depends solely on the number of ‘active cells per pattern’ and it is completely independent of how many patterns are stored in the network. When just 5 ‘active cells per pattern’ are used to represent a memory, then the recall performance is best for all three models across all overlaps and irrespective of number of stored patterns. When 10 ‘active cells per pattern’ are used to represent a memory, the performance of all three models are comparably similar when 5 or 10 patterns are stored and across overlap percentages. When 20 ‘active cells per pattern’ are used to represent a memory, then even for just 5 patterns stored, the recall performance is consistently best for ‘model 1’ and consistently worst for ‘model 2’ across all overlaps. ‘Model 3’ performance is between ‘model 1’ and ‘model 2’. The performances of ‘model 2’ and ‘model 3’ get worse as overlap increases (from 10% to 40%).

Figures 4 and 5 compare and contrast the recall performance of models 1, 2, and 3 against number of ‘active cells per pattern’ for various overlaps (0%, 10%, 20% and 40%) and stored patterns (5 or 10). When 5 or 10 ‘active cells per pattern’ are used to represent a memory, then the recall performances of all three models when number of stored patterns were 5 or 10 was exactly the same at 0%, 10%, 20% and 40%, respectively. This means that the number of patterns stored in the network did not affect its recall quality. When ‘active cells per pattern’ were increased (from 10 to 20), then the recall qualities of models 2 and 3 progressively got worse as overlap between patterns increased (from 0% to 40%). ‘Model 1’ recall quality was consistently best ($C = 1$) across ‘active cells per pattern’, stored patterns, and overlap conditions.

Why was ‘model 1’ performance so consistently better than ‘model 2’ and ‘model 3’ across all conditions? Why the recall quality of ‘model 1’ was always perfect ($C = 1$) even when more patterns were stored in the network, more/less ‘active cells per pattern’ were used to represent a memory and greater percentages of overlap between patterns were used? As we stated in section “2.7 – Model selection”, ‘model 1’ was the model where CA3 feedforward excitation to BSC was progressively increased as more and more patterns were stored, while the BSC inhibitory effect to PC dendrites was held fixed. ‘Model 3’ was the model where PC feedback excitation to BSC was progressively increased as more and more patterns were stored, while the BSC inhibitory effect to PC dendrites was held fixed. ‘Model 2’ was the model where the exact opposite took place: the inhibitory effect of BSC to PC dendrites progressively increased as more and more patterns were stored in the network, while keeping the BSC firing rate constant. In all simulations, ‘model 1’ outperformed ‘model 3’ across all conditions (overlaps and ‘active pattern cells’). This was due to the fact that in ‘model 1’ BSC was excited by 100 CA3-PCs at high frequency (40 Hz), whereas in ‘model 3’ BSC was excited by 20 CA1-PCs that fired once or twice. Since in ‘model 1’ the BSC firing frequency response is higher than in ‘model 3’, then the postsynaptic effect of BSC on the PC dendrites in ‘model 1’ is higher in frequency and duration (but not in amplitude) than in ‘model 3’ (see Fig. 6A & 6C). Thus, ‘model 1’ has a better success at removing spurious activities and improving recall quality than ‘model 3’. Since the BSC frequency response in ‘model 2’ was fixed, but its postsynaptic effect (weight) on PC dendrites increased, then the amplitude of the inhibitory postsynaptic potentials (IPSPs) on PC dendrites increased (compared to the IPSP amplitudes in

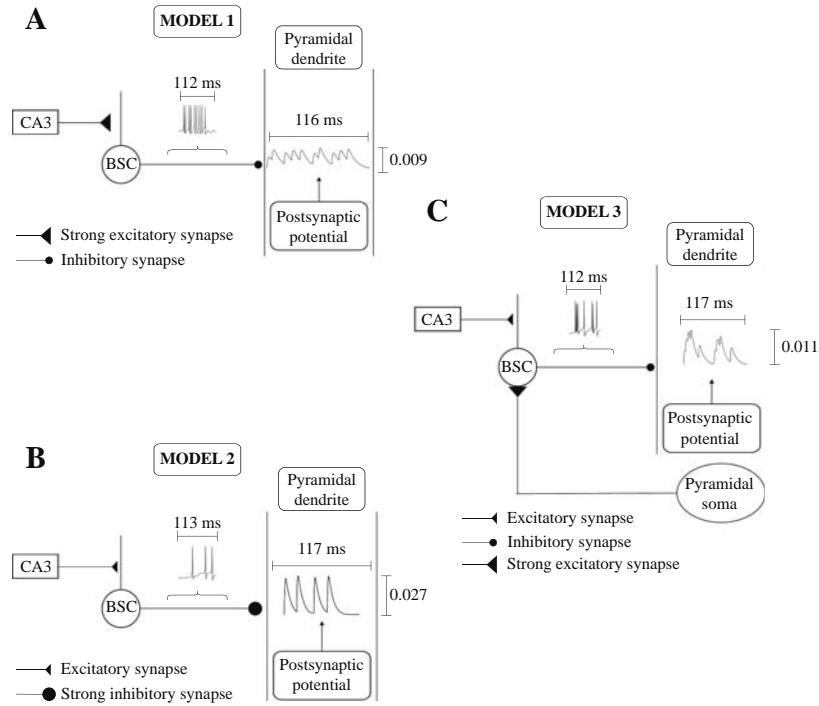


Fig. 6. Schematic drawing of presynaptic BSC firing response and inhibitory postsynaptic potentials (IPSPs) on PC dendrites in (A) ‘model 1’, (B) ‘model 2’ and (C) ‘model 3’.

models 1 and 3), but their frequency response was low (lower than in models 1 and 3; see Fig. 6B). Each IPSP decayed to almost zero before another IPSP was generated post-synaptically on PC dendrites.

4 Conclusions

A bio-inspired neural model of mammalian hippocampal CA1 region [5] was employed to systematically evaluate its mean recall quality against number of stored patterns, percent overlaps and ‘active cells per pattern’. We modulated the strength of selective excitatory and inhibitory pathways to BSC and PC dendrites as more and more patterns were stored in the network of 100 CA1-PCs and this resulted into three models, the performances of which were compared against each other. Model 1 recall performance was excellent ($C = 1$) across all conditions. Model 2 performance was the worst. A key finding of our study is that the number of ‘active cells per pattern’ has a massive effect on the recall quality of the network regardless of how many pattern are stored in it. As the number of dedicated cells representing a memory (‘active cells per pattern’) decrease, the memory capacity of the CA1-PC network increases, so interference effects between stored patterns decrease, and mean recall quality increases. Another key finding of our study is that increased firing frequency response of a presynaptic inhibitory cell (BSC) inhibiting a network of PCs has a better success at removing spurious activity at the network level and thus improving recall quality than an increased synaptic efficacy of a presynaptic inhibitory cell (BSC) on a postsynaptic PC while keeping its presynaptic firing rate fixed.

Acknowledgements. This work was supported in part by EU Horizon 2020 through Project ULTRACEPT under Grant 778062.

References

1. Suzuki, W.A.: Making new memories: the role of the hippocampus in new associative learning. *Ann. N. Y. Acad. Sci.* **1097**, 1–11 (2007)
2. Eichenbaum, H., Dunchenko, P., Wood, E., Shapiro, M., Tanila, H.: The hippocampus, memory and place cells: is it spatial memory or a memory of space? *Neuron* **23**, 209–226 (1999)
3. Cutsuridis, V., Graham, B.P., Cobb, S., Vida, I.: *Hippocampal Microcircuits: A computational modeller’s resource book*, 2nd edn. Springer, Cham (2019). <https://doi.org/10.1007/978-3-319-99103-0>
4. Treves, A., Rolls, E.: Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* **2**, 189–200 (1992)
5. Cutsuridis, V., Cobb, S., Graham, B.P.: Encoding and retrieval in a model of the hippocampal CA1 microcircuit. *Hippocampus* **20**, 423–446 (2010)
6. Klausberger, T., et al.: Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* **421**, 844–848 (2003)
7. Klausberger, T., Marton, L.F., Baude, A., Roberts, J.D., Magill, P.J., Somogyi, P.: Spike timing of dendrite-targeting bistratified cells during hippocampal network oscillations in vivo. *Nat. Neurosci.* **7**, 41–47 (2004)
8. Poirazi, P., Brannon, T., Mel, B.W.: Arithmetic of subthreshold synaptic summation in a model of CA1 pyramidal cell. *Neuron* **37**, 977–987 (2003)
9. Poirazi, P., Brannon, T., Mel, B.W.: Pyramidal neuron as a 2-layer neural network. *Neuron* **37**, 989–999 (2003)
10. Santhakumar, V., Aradi, I., Soltetz, I.: Role of mossy fiber sprouting and mossy cell loss in hyperexcitability: a network model of the dentate gyrus incorporating cell types and axonal topography. *J. Neurophysiol.* **93**, 437–453 (2005)
11. Buhl, E.H., Szilágyi, T., Halasy, K., Somogyi, P.: Physiological properties of anatomically identified basket and bistratified cells in the CA1 area of the rat hippocampus in vitro. *Hippocampus* **6**(3), 294–305 (1996)
12. Buhl, E.H., Han, Z.S., Lorinczi, Z., Stezhka, V.V., Kapnup, S.V., Somogyi, P.: Physiological properties of anatomically identified axo-axonic cells in the rat hippocampus. *J. Neurophysiol.* **71**(4), 1289–1307 (1994)
13. Cutsuridis, V.: Improving the recall performance of a brain mimetic microcircuit model. *Cogn. Comput.* **11**, 644–655 (2019). <https://doi.org/10.1007/s12559-019-09658-8>
14. Hines, M.L., Carnevale, T.: The NEURON simulation environment. *Neural Comput.* **9**, 1179–1209 (1997)
15. Borhegyi, Z., Varga, V., Szilagy, N., Fabo, D., Freund, T.F.: Phase segregation of medial septal GABAergic neurons during hippocampal theta activity. *J. Neurosci.* **24**, 8470–8479 (2004)
16. Petersen, C.C.H., Malenka, R.C., Nicoll, R.A., Hopfield, J.J.: All-or none potentiation at CA3-CA1 synapses. *Proc. Natl. Acad. Sci. USA* **95**, 4732–4737 (1998)