

Model of the Hippocampal Learning of Spatio-temporal Sequences

Julien Hirel, Philippe Gaussier, and Mathias Quoy

Neurocybernetic team, ETIS, CNRS - ENSEA - University of Cergy-Pontoise, F95000

Final draft version, accepted for publication as:

Hirel, J., Gaussier, P., Quoy, M.: Model of the hippocampal learning of spatio-temporal sequences. In Diamantaras, K., Duch, W., Iliadis, L., eds.: Artificial Neural Networks – ICANN 2010. Volume 6354 of Lecture Notes in Computer Science., Springer Berlin / Heidelberg (2010) 345–351 10.1007/978-3-642-15825-4_46.

The original publication is available at www.springerlink.com

Abstract. We propose a model of the hippocampus aimed at learning the timed association between subsequent sensory events. The properties of the neural network allow it to learn and predict the evolution of continuous rate-coded signals as well as the occurrence of transitory events, using both spatial and non-spatial information. The system is able to provide predictions based on the time trace of past sensory events. Performance of the neural network in the precise temporal learning of spatial and non-spatial signals is tested in a simulated experiment. The ability of the hippocampus proper to predict the occurrence of upcoming spatio-temporal events could play a crucial role in the carrying out of tasks requiring accurate time estimation and spatial localization.

1 Introduction

When an animal moves along a particular trajectory, its movement can be interpreted with two complementary perspectives: (i) as a purely spatial strategy using visual input to provide information about the current localization and adapt the behavior accordingly; (ii) as a series of actions cued to particular events, triggered by an internal time representation or path integration information. Lagarde et al. [1] showed that a robot can learn a desired trajectory both by using place-action associations or timed sequences of actions. But what about the tasks where both spatial and temporal dynamics are needed ? Is temporal and spatial information processed in different structures and somehow integrated in a larger neural network allowing spatio-temporal strategies, or is there a neural substrate directly capable of learning both spatial and temporal properties ? Some experimental tasks involve both spatial and temporal aspects [2]. In that context the animal needs time estimation, reward prediction and navigation capabilities, and a way to merge these modalities.

In this paper we propose a model where the hippocampus can learn and predict transitions between multi-modal sensory events. The model provides an internal time representation allowing the learning of timed sequences of stimuli. Using a memory of past events, we can then predict the occurrence of future transitory events or estimate the evolution of a continuous signal. Using the same architecture we are able to learn the temporal properties of various signals, spatial and/or non-spatial, in accordance to the multi-modality observed in hippocampal cells [3]. The model is tested in a simulation where a single neural network is used to learn to predict place cell activities and an unrelated non-spatial sequence of sensory stimuli.

2 Model

We previously presented a model of the hippocampus as an associative memory capable of learning the correlation between past and present sensory states and of predicting accessible states [4]. This architecture has been mainly used in two experimental contexts. First, a version using a simple memory of past states has been used in robot navigation to learn transitions between places [5]. Second, a timed memory inspired from the spectral timing model [6] was used to learn timed sequences of sensory events [7] and later for the reproduction of sequences of motor actions in a robotic experiment [1]. As previously discussed, the carrying out of certain tasks requires a strategy which integrates both temporal and spatial components, the two closely intertwined. To process all sorts of information, spatial or not, we gave our neural network the ability to predict the occurrence of punctual events as well as estimate the evolution of continuous signals over time. As an evolution over previous models, the “one-shot” learning of timings was replaced by a continuous learning. This allows the predictions to adapt their timing over time or to be forgotten if they are never fulfilled.




Fig. 1. Model of the state prediction neural network.

In our model, EC contains the current sensory states of the robot. Multimodal signals (e.g. vision, sound, odometry etc.) are integrated. A *Winner-Take-All* (WTA) competition ensures that the activity of the most active state (i.e. current state) is transmitted to the DG memory. Each state is connected to a

corresponding battery of neurons with various temporal properties in DG [7]. When the current state changes, the memory is reset and the battery corresponding to the new state is activated. This battery then gives a time trace of the delay elapsed since the state was triggered. Topological distal connections transmit state activity from EC to CA3. The activity of CA3 neurons in (1) is computed using DG-CA3 synapses only and corresponds to the prediction of future states. The learning (2) takes place in the DG-CA3 synaptic weights and is based on the Normalized Least Mean Square (NLMS) algorithm [8], with the addition of a synaptic learning modulation η .

$$x_i^{CA3}(t) = \sum_{k \in DG} W_{ik}^{DG-CA3} \cdot x_k^{DG} - \theta \quad (1)$$

$$W_{ij}^{DG-CA3}(t + dt) = W_{ij}^{DG-CA3}(t) + \alpha \cdot \eta_i(t) \cdot \frac{(x^{EC}(t)_i - x_i^{CA3}(t))}{\sum_{k \in DG} x_k^{DG}(t)^2 + \sigma_1} \cdot x_j^{DG}(t) \quad (2)$$

where θ is an activity threshold, W_{ij}^{DG-CA3} is the synaptic weight from DG neuron j to CA3 neuron i , α is the learning rate, η_i a learning modulation (see eq. 3), x_i^{EC} is the activity transmitted to neuron i by EC (i.e. the target signal for the LMS). The particularity of the NLMS over the standard LMS is the normalization term $\sum_k x_k^{DG}(t)^2 + \sigma_1$ representing the energy of the input vector. σ_1 is a small value used to avoid the divergence of the synaptic weights for very low DG values.

Each neuron of a DG battery displays a pattern of activity defined by a Gaussian function of the time elapsed since the battery was activated, with its own particular mean and variance [7]. Using the NLMS rule and a continuous signal as EC input, the neural network is capable of predicting the evolution of the signal using its internal time representation. When working with transitory input signals, like punctual events generating short bursts of neural spiking, the model learns to predict the timing of the event in relation to the time passed since the last event. The result is a bell shaped rate-coded activity, with a peak predicting the event. Since learning is performed online, the use of a LMS algorithm is problematic (the samples are not randomly distributed). As soon as a battery of DG cells is activated, the NLMS starts to learn to approximate the EC input signal and minimize the estimation error. The prediction of an event can start well before the onset of the burst of neural activity coding for this event. As a result, a predictive output signal is present while there is no activity on EC. The NLMS learning rule will converge to correct this error, thus leading to the decay of synaptic weights. Once the predicted event occurs and neural activity arises in EC, the previous decay will be balanced by the new learning of the input signal. Yet, the NLMS learning rule intrinsically leads to the same decay and learning speeds. Since transitory signals have short periods of activity (learning phase) and their prediction activity can start well in advance of their occurrence (decay phase), the previous learning of a transition is quickly forgotten. In order to learn online transitory events, the CA3 learning rule needs to be more reactive to the short period of EC activity than to the periods with no input activity. Therefore, our model includes a local modulation η of the synaptic learning rule. Its equation (3) uses the difference between the current value of the

EC input signal and its mean value over a certain period of time. This results in the modulation of the synaptic learning of DG-CA3 plastic proximal connections by the variability of the activity coming from EC-CA3 distal connections. The modulation leads to a higher sensitivity of the associative learning to transitory signals: rapidly changing signals are learned faster than stable ones.

$$\eta_i(t) = |x_i^{EC}(t) - m_i^{EC}(t)| + \sigma_2 \text{ with } m_i^{EC}(t) = \gamma \cdot m_i^{EC}(t-1) + (1-\gamma) \cdot x_i^{EC}(t) \quad (3)$$

where m_i^{EC} is a sliding mean of x_i^{EC} , σ_2 is a low value setting a minimal learning modulation for CA3 and γ a parameter controlling the balance between past and current activities in the computation of the sliding mean.

Interestingly, the model is both able to approximate continuous input signals and to predict transitory events. Due to the modulation η , the synaptic weights are slowly decaying when the predicted transitions do not occur. As a side effect, the system is slower in the learning of stable continuous signals.

3 Results and Discussion

A navigation experiment where a simulated robot follows a defined trajectory was conducted (Fig. 2). Place cells are learned on a minimum activity threshold basis and their activity is computed using the azimuth of the recognized landmarks in the visual field. Each battery of DG cells is composed of 15 neurons giving a spectral decomposition of time on a period of 5 seconds.




Fig. 2. 9m “∞-loop” trajectories taken by the agent in the 3x3m simulated environment. 20 perfectly identifiable landmarks are placed regularly along the walls. The robot has a constant speed of 0.5m/s. One loop requires at least 18s to be completed. Time is discretized into a series of 100ms steps. The colors correspond to the most activated place cell at each location (labeled by letters).

Three types of EC signals are presented concurrently as input to CA3:

Raw place cell activity: Activity of all place cells. The DG memory gives a trace of the time passed since the arrival on the last place. The neural network learns to predict the evolution of place cell activity based on the time spent in the current place.

“Place entered” transitory signal: When the current place changes, EC triggers a short burst of activity. The DG memory is the same as for raw place cell activity. The timing of transitions from the current place to accessible neighbor places is predicted.

Non-spatial timed sequence of events: A recurring series of events, signaled by short transitory activity, forms a sequence which could represent the interaction with a person presenting various stimuli. A separate WTA competition and DG memory keeps a trace of the last event. The prediction of the timed sequence of events is learned by CA3, independently of the spatial context.

After a long period (over 50 ∞ -loops) of learning¹, we turn off synaptic modifications in order to analyze the predictions of the CA3 neurons based solely on DG activity. Figure 3 shows the EC input signal and CA3 predictions after the learning, over a period corresponding to one complete ∞ -loop. Even though the three types of signal are learned online concurrently by the same neural network, they are represented separately for more clarity. We can see that the system correctly learns to predict the evolution of place cell activity, the timing of place transitions and the sequence of non-spatial events. Small drops of raw place cell prediction activity can be observed when the most active place cell changes. This is due to the reset of the DG memory and the short interval of time (100ms) needed by the first DG cells to be activated. The sum of the mean square errors (MSE) of raw place cell activity prediction, after learning and over a period of 10 ∞ -loops, is 0.035. Event predictions are learned for both spatial information (a change of the most active place cell) and non-spatial information (a repeated sequence of events, i.e. 1-2-3-2-1). All the possible events are predicted by a bell-shaped pattern of activity with the maximum value corresponding to the expected timing arrival. Predictions with a low activity can be observed for spatial events. The noise on the place cell activity can indeed lead to rare transitions between non-successive places, the low activity is a characteristic of the rareness of these occasions.




Fig. 3. **a)** Continuous prediction of all place cell activity based on the time passed since entry on the current place. **b)** Event prediction of the timing of the arrival in the next place. **c)** Prediction of the timing of the occurrence of the next non-spatial event.

In order for the network to learn properly, the mean in (3) needs to use a short sliding window ($\gamma = 0.5$ in the simulation) so that, when the prediction of an event begins, the previous peak of activity for this event is already forgotten. This supposes local synaptic learning mechanisms for pyramidal neurons in CA3 with properties of short-term memory. A learning rate $\alpha = 0.5$ can seem high for LMS-based learning. Yet, it allows the fast learning of transitory events (which

¹ Parameters used for learning: $\alpha = 0.5$, $\sigma_1 = 0.01$, $\theta = 0.05$, $\sigma_2 = 0.01$, $\gamma = 0.5$.

have a high η value during the short period of activity) and the slow learning of continuous signals (which have a low η value due to their low variability) needed for the stability of the LMS algorithm. Finally, an interesting property of the neural network is that, as the network learns the timing over and over, the peak of event predictions becomes more and more narrow. As fig. 3 **b)** and **c)** show, the long period of learning in the experiment leads to very sharp predictions, with a peak of activity centered on the predicted occurrence of the next event.

In conclusion, we have presented a neural network capable of learning spatial and non-spatial, continuous and transitory signals. The learned association between past and current signals results in predictive capabilities. Experimental observations of our model suggest that the learning of the timing goes through two phases: i) a short learning phase during which the amplitude of the predictions rises to reach a maximal value; ii) A long adaptation phase during which the prediction peak of activity gets progressively narrower and centered on the precise expected timing of the next event. In repetitive tasks where place transitions or sensory events have low variability in their timings, we expect to observe progressively sharper peaks of activity and reduced temporal precession for predictions. As the learning starts, it is not always clear which aspects of the task will be crucial to perform efficiently. A hypothesis is that dedicated structures learn specific components of the task (e.g. cerebellum for timed conditioning, basal ganglia for reward expectations, prefrontal cortex for movement inhibition etc.) and that superfluous information is discarded. We propose a model of the hippocampus as an associative memory extracting not only spatial but also temporal characteristics of the sensory stimuli experienced during the task. This allows to reconcile Eichenbaum's view [9] of the hippocampus as a memory and O'Keefe's view [10] of the hippocampus as a cognitive map. Finally the ability of the neural network to predict the evolution of state activity could be used in pair with the actual state activity to compute the error between the prediction and the current context. In well rehearsed tasks where the prediction should have reached a certain level of accuracy, the error signal could be used to detect anomalies or novelty. The robot would then be able to detect that the conditions of the task may have changed and adapt its behavior accordingly.

Acknowledgments This work is supported by the CNRS, as part of a PEPS project on neuroinformatics, the DGA and the IUF. We thank JP. Banquet, B. Poucet and E. Save for useful discussions.

References

1. Lagarde, M., Andry, P., Gaussier, P., Giovannangeli, C.: Learning new behaviors : Toward a control architecture merging spatial and temporal modalities. In: Workshop on Interactive Robot Learning - International Conference on Robotics: Science and Systems (RSS 2008). (June 2008)
2. Hok, V., Lenck-Santini, P.P., Roux, S., Save, E., Muller, R.U., Poucet, B.: Goal-related activity in hippocampal place cells. *J Neurosci* **27**(3) (Jan 2007) 472–482

3. Wiener, S.I., Paul, C.A., Eichenbaum, H.: Spatial and behavioral correlates of hippocampal neuronal activity. *J Neurosci* **9**(8) (Aug 1989) 2737–2763
4. Banquet, J.P., Dreher, J.C., Revel, A., Gunther, W.: Space-time, order and hierarchy in fronto-hippocampal system: A neural basis of personality. In: *Cognitive Science Perspectives on Personality and Emotion*, Elsevier (1997) 123–189
5. Cuperlier, N., Quoy, M., Gaussier, P.: Neurobiologically inspired mobile robot navigation and planning. *Front Neurobotics* **1** (2007) 3
6. Grossberg, S., Schmajuk, N.A.: Neural dynamics of adaptive timing temporal discrimination during associative learning. *Neural Netw.* **2**(2) (1989) 79–102
7. Moga, S., Gaussier, P., Banquet, J.P.: Sequence learning using the neural coding. In: *IWANN'03: Proceedings of the 7th International Work-Conference on Artificial and Natural Neural Networks*, Berlin, Heidelberg, Springer-Verlag (2003) 198–205
8. Nagumo, J.: A learning method for system identification. *IEEE Trans. Autom. Control* **12**(3) (1967) 282–287 NLMS reference.
9. Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., Tanila, H.: The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron* **23**(2) (Jun 1999) 209–226
10. O'Keefe, J., Nadel, L.: *The hippocampus as a cognitive map*. Oxford University Press (1978)