Neural Networks 62 (2015) 39–46

Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2015 Special Issue

Neural coordination can be enhanced by occasional interruption of normal firing patterns: A self-optimizing spiking neural network model

Alexander Woodward [a,*], Tom Froese [b,c], Takashi Ikegami [a]

[a] Graduate School of Arts and Sciences, The University of Tokyo, Komaba, Tokyo 153-8902, Japan

[b] Departamento de Ciencias de la Computacion, Instituto de Investigaciones en Matematicas Aplicadas y en Sistemas, Universidad Nacional Autonoma de Mexico, Mexico

[c] Centro de Ciencias de la Complejidad, Universidad Nacional Autonoma de Mexico, Mexico

ARTICLE INFO

Article history:
Available online 16 September 2014

Keywords:
Self-optimization
Hopfield network
Spiking neurons
Global neural coordination
Psychedelics
Altered states of consciousness

ABSTRACT

The state space of a conventional Hopfield network typically exhibits many different attractors of which only a small subset satisfies constraints between neurons in a globally optimal fashion. It has recently been demonstrated that combining Hebbian learning with occasional alterations of normal neural states avoids this problem by means of self-organized enlargement of the best basins of attraction. However, so far it is not clear to what extent this process of self-optimization is also operative in real brains. Here we demonstrate that it can be transferred to more biologically plausible neural networks by implementing a self-optimizing spiking neural network model. In addition, by using this spiking neural network to emulate a Hopfield network with Hebbian learning, we attempt to make a connection between rate-based and temporal coding based neural systems. Although further work is required to make this model more realistic, it already suggests that the efficacy of the self-optimizing process is independent from the simplifying assumptions of a conventional Hopfield network. We also discuss natural and cultural processes that could be responsible for occasional alteration of neural firing patterns in actual brains.

1. Introduction

The class of recurrent Hopfield neural networks, first described by Hopfield (1982), has traditionally been employed for two distinct kinds of tasks. On the one hand, these networks can be trained to form an associative memory of neural activity patterns. Neural activity is set to the activation pattern to be memorized and Hebbian learning is applied in order to turn that pattern into an attractor. One drawback is that spurious attractors are easily formed and these do not represent any target pattern. On the other hand, Hopfield networks can also be used to find solutions to constraint satisfaction problems (Hopfield & Tank, 1985). The connection weights are set to represent the constraints between the components of the target problem, the network's activity is initialized to some starting configuration, and the activity is then allowed to converge to an attractor, which at the same time represents a possible solution to the problem represented by the weights. This process of convergence can be understood as a coordination of component activity so as to satisfy the most constraints given the starting configuration. However, as is well known, the state space of a complex Hopfield network typically exhibits many different attractors of which only a small subset is globally optimal; the rest are local optima that fail to take full advantage of the possibilities of coordination.

Watson, Buckley, and Mills (2011); Watson, Mills, and Buckley (2011a) recently discovered that Hopfield networks that combine these two tasks manage to overcome both types of drawbacks. They modified the standard constraint satisfaction procedure by making it iterative and including Hebbian learning. As per usual, the weights of the network are set to represent a specific constraint satisfaction problem. Then they used the following itinerant routine: (1) the neural network activity is initialized to a random configuration, (2) the activity is allowed to converge to an attractor, and (3) after this point a small amount of Hebbian learning is applied. What is the effect of repeating this three-step procedure? As might be expected, the neural network forms an associative memory. However, in this case it is not a memory of external patterns, but rather of the different attractor configurations that

* Corresponding author. Tel.: +81 03 5465 8853.
E-mail addresses: alex.w.nz@gmail.com, alex@sacral.c.u-tokyo.ac.jp (A. Woodward).

http://dx.doi.org/10.1016/j.neunet.2014.08.011
0893-6080/

the network has visited. Over time the network will thereby reconfigure its weight space until most (if not all) initial activity configurations lead to the same attractor, which happens to be one of the best solutions to the original constraint satisfaction problem.

Two properties of Hopfield neural networks are responsible for this useful self-optimization process. First, it has been proven that there is a positive correlation between the width (localizability) and depth (optimality) of a basin of attraction (Kryzhanovsky & Kryzhanovsky, 2008), which means that better constraint solutions are visited comparatively more frequently and are therefore reinforced more often. Second, the self-optimization process takes advantage of the learning neural network's ability to generalize over the training set, i.e. the visited attractors. In this case the reinforcement of spurious attractors is actually desirable. For as long as the problem in weight space is decomposable in some manner, the reinforcement of a visited attractor at the same time reinforces other attractors that are partially composed of similar configurations—even if the network has not previously encountered them after one of the re-initializations. In this manner the basins of attraction of still unvisited global optima will become enlarged, and therefore more easily found, even if they are normally extremely difficult to locate. This is effectively the same as if the neural network is translating its original constraint optimization problem into a higher-dimensional organizational space to make it easier to solve, but without making use of any a priori knowledge of the problem domain (Watson, Mills, & Buckley, 2011b).

Given the generality of the mathematics underlying the Hopfield network, which is isomorphic to the famous class of Ising models in statistical mechanics (Rojas, 1996, Ch. 13), as well as the simplicity of the self-optimization process, we can expect this process to govern the emergence of coordination in a wide range of systems. Even the need for true Hebbian learning can be relaxed. In the case of social systems it has been shown that habituation of the behaviors that constitute attractor configurations is sufficient to realize a similar structural self-optimization process (Davies, Watson, Mills, Buckley, & Noble, 2011). Nevertheless, it remains to be verified that the process proposed by Watson and colleagues remains effective when it is implemented in more biologically realistic neural networks, and to provide an interpretation of the necessary periodic deviations from converged behavior (i.e., step 1).

In this work we created a spiking neural network model that emulates the properties of a traditional Hopfield network with saturated linear transfer (rather than binary threshold) functions, and with real-valued (rather than integer) weights. Our main aim was to demonstrate that the combination of Hebbian learning with occasional alteration of normal neural network activity also leads to the emergence of global neural coordination in such a spiking neural network. Although further modeling work is required to confirm that this self-optimizing process can be operative in even more realistic neural networks, here we managed to show that it is independent of the simplifying assumptions of the conventional Hopfield network. We only interpret the spiking network as a Hopfield network in order to demonstrate that self-optimization is indeed taking place in an equivalent manner.

Additionally, Hebbian learning is a rate-based learning method and we have proposed a form of Hebbian learning in a timing based system, by using heterosynaptic plasticity (Bailey, Giustetto, Huang, Hawkins, & Kandel, 2000; Huang, Pittenger, & Kandel, 2004) and spike-timing dependent plasticity (STDP), equivalent to that found in traditional Hopfield networks, thus making a connection between rate-based and temporal neural encoding systems. Temporal coding based systems have a number of advantages such as not being affected by synaptic depression and being able to achieve a high rate of computation at biological realistic firing rates (Maass & Natschlaeger, 1997). In the same paper, Mass et al. have pointed out that from recent experiments 'it is in fact questionable whether biological neural systems are able to carry out analog computation with analog variables represented as firing rates' (p. 355).

The rest of this article unfolds as follows. First, we discuss our methods in more detail, paying special attention to how we applied the ideas gained from studies with Hopfield neural networks to the spiking neural network model. Second, we present the results of our investigation, which demonstrate that Hopfield dynamics and the process of self-optimization of neural coordination identified by Watson and colleagues is also effective in spiking neural networks. Finally, we evaluate the plausibility of the spiking neural network model when compared to real nervous systems. We also briefly discuss what could be the natural and cultural causes of occasional alteration of normal neural activity in human brains.

2. Methods

The large number of recurrent connections in the brain seems to be the mechanism behind its associative memory capabilities. A well known computational neural network architecture that also exhibits such properties is the Hopfield network, first described in Hopfield (1982), and with graded neuron response in Hopfield (1984). In the following we show how it is possible to transfer some of the key concepts of the Hopfield network to a biologically more realistic spiking neural network.

In early versions of the Hopfield network (Hopfield, 1982) the binary state of a neuron was taken to abstractly represent whether that neuron was not firing (0) or firing at maximum rate (1). In later versions (Hopfield & Tank, 1985), as well as in more recent elaborations of this kind of network architecture such as including negative self-feedback (Nozawa, 1992), or the continuous-time recurrent neural network (Beer, 1995), a nonlinear function of a neuron's state is typically interpreted to be its mean firing rate.

It is therefore reasonable to assume that, conversely, the firing rate of a spiking neuron can be interpreted as the state of a traditional Hopfield neuron. However, as alluded to in the introduction, doubt has been raised as to whether firing rates in a biological neural system would be even appropriate for carrying out Hopfield network dynamics. Since temporal coding systems have the potential to achieve a high rate of computation at biological realistic firing rates we therefore describe a way to implement Hopfield network dynamics using temporal coding in a spiking neural network.

To measure constraint satisfaction by means of neural coordination, the spike-timing encoding of analog values can then be inserted into the Hopfield energy function where attractor states of the network correspond to local minima in the energy landscape (by convention, larger negative values in the Hopfield energy function are more optimal). Our proposed heterosynaptic Hebbian learning approach, described in Section 2.3.3, can then be used to update weights in the spiking network as part of the self-optimization process.

2.1. The Hopfield network

A Hopfield network is a fully interconnected neural network, usually with symmetric connections between nodes, where H represents the network state: $H = \langle s_1, \dots, s_n \rangle \in [0, 1]^n$.

Hopfield dynamics can be described through updates to neuron states; for the ith neuron s_i:

$$s_i(t+1) = \theta \left[\sum_{j}^{N} \omega_{ij} s_j(t) \right] \tag{1}$$

where ω_{ij} is the weight between neurons i and j, collectively described by $\Omega = \langle \omega_1, \dots, \omega_n \rangle \in [-1, 1]^n$ and θ is a transfer