Neural Information Processing - Letters and Reviews Vol.7, No.2, May 2005

LETTER

Coherence, Memory and Conditioning : A Modern Viewpoint

Raul C. Mureşan

Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe University
Max-von-Laue-Strasse 1, 60438 Frankfurt am Main

Center for Cognitive and Neural Studies (Coneural)
Str. Saturn 24, 400504 Cluj-Napoca, Romania

Faculty of Automation and Computers, Technical University of Cluj Napoca
26-28 Baritiu Street, 400027 Cluj-Napoca, Romania
E-mail: raulmuresan@yahoo.com

Gordon Pipa

Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe University
Max-von-Laue-Strasse 1, 60438 Frankfurt am Main

Răzvan V. Florian

Center for Cognitive and Neural Studies (Coneural)
Str. Saturn 24, 400504 Cluj-Napoca, Romania
E-mail: florian@coneural.org

Diek W. Wheeler

Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe University
Max-von-Laue-Strasse 1, 60438 Frankfurt am Main

(Submitted on March 16, 2004)

Abstract— We present a novel viewpoint on the complexity of neural mechanisms, addressing some aspects of cortical processing, like memory, context modulation and coherence. Our simulation studies show how relatively small, recurrent microcircuits can interact with populations of neurons, achieving spontaneous memory / recall effects. Here, context modulation can play an important part in the conditioned recall. The emphasis is on the dynamical self-sustained activity of the recurrent microcircuit viewed as a "mental state" which can be interpreted by a target output population and controlled by external waves of activity (mainly external inhibitory waves). We also consider the problem of self-organization and dynamical interaction between neuronal populations, which could have a key role in inference, memory and behavior.

Keywords— Neural microcircuit, coherence, memory, conditioning, context

1. Introduction

An obvious aspect of information processing in the brain is the fact that every sensorial input is assimilated over an already existing internal "mental state" [1]. There is also evidence that most of the neural microcircuits are recurrent and that connectivity between cortical neurons is really high (on the order of 10^3 - 10^4) [2]. Moreover, it has been shown that highly irregular spike trains can be obtained with fixed random connectivity networks, questioning the source of the neural noise [3,4,5]. Within this context, we might attribute, at least partially, the apparently noisy, spontaneous neural activity to the "mental state" of a recurrent microcircuit (the so-called network effect).

It seems essential to analyze the neural activity of real neurons from a new computational perspective, which can reveal the dynamical nature of processing. Such an approach is really hard, since we deal with enormously large dynamical systems. Nonetheless, extensive research is already emerging, a good example being the work of Maass, Natschläger and Markram [6]. It has been emphasized that for large systems like neural microcircuits an exact link between the function of an individual neuron and its activity cannot always be found with traditional methods, its dynamics and interactions being far too complex for a comprehensive analysis. An alternative approach is to look at the activity of a population of neurons rather than to single units. This in turn, has its own limitations since one has to determine the target population (not always an easy task) and apply statistical analysis which might sometimes reveal irrelevant functional aspects.

An emergent new technique of studying such large, complex, dynamical neural systems is to build and test large scale models, with biologically plausible components [7,8]. Using this technique, a novel approach to understanding the activity of high dimensional recurrent microcircuits has been used in the "liquid state machine" (LSM) [6] and "the echo state approach" [9]. In both cases, a population of readout neurons is trained (by linear regression for example) to observe the highly dynamic processing in a recurrent module and convert the unstable dynamics into stable, reliable output activity.

Although, the recurrent microcircuit is highly biologically plausible, the existence of separate, independent readout populations [6] in the brain, that can be trained by regression methods, seems implausible. More likely, there is a continuum of microcircuits interconnected in a complex, evolution guided manner, where parts of a recurrent microcircuit are readout for other microcircuit modules and vice-versa. Plausible mechanisms like spike-timing dependent plasticity (STDP) could be the basis for the interoperability of different spatial modules of the continuous microcircuit sheet. Development could be guided by behavioral experience, by the interaction between the environment and the organism [10, 11, 12], in addition to genetic constraints. Dynamic associations might arise while the brain could self-wire itself in a complex way, perhaps too complex for a comprehensive analysis. Within this framework, we questioned the existence of a strong, unifying principle that guides cortical processing, called "The Coherence Principle" [13].

2. Coherence and Microcircuits

Coherence is a universal yet simple principle, which states that a given interpretation is only "sound" in a given context. Thus, we can never separate the interpretation from the context. Also, an important observation is the fact that neural activity is most of the time causal, i.e. the activity of each neuron depends on the context of afferent activity. Moreover, we have to take into account the increasing complexity of neural activity, as one proceeds upstream in a sensorial processing hierarchy in the brain. While the primary layers seem to have deterministic responses, mainly driven by the input (as is the case in the primary visual cortex), deeper processing structures have a complex activity with complex firing patterns of the neurons. As one proceeds to higher levels of sensorial processing, the complexity of interactions between neural areas grows significantly. We could then separate the afferent populations of a cell into "input populations" and "context populations" (Figure 1). Clearly, this separation depends on the "interpretation" derived from the cell's activity. There are many possible ways for segregating these two classes. A simple and maybe plausible segregation would be to label as input populations all the afferent neurons that are placed "before" the cell in the sensorial processing chain. Then, context populations would be all the other neurons that contribute to the activity of the target neuron. Lateral interaction and feedback neurons belong, in this case, to the context population. This definition can be extended from single neuron targets to target populations of neurons.

The same input could trigger a whole different range of interpretations for different context activities. The activity of a cell or population that has important afferent activity, different from the input, cannot be assumed to

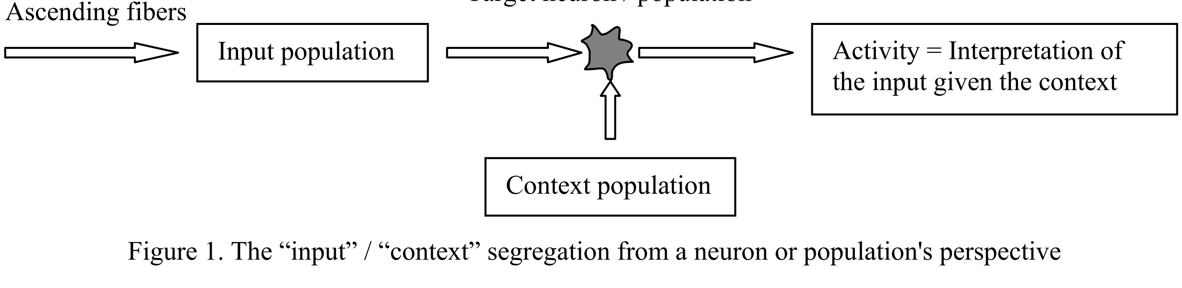

Figure 1. The “input” / “context” segregation from a neuron or population's perspective

be simply a function of the input. This is why a transfer function to characterize the activity in higher visual cortical areas, such as V4, based on the input stimulus, is very hard if not impossible to define.

Microcircuit populations might cooperate to encode complex contexts and modulate other populations to generate interpretations. Competition, which has been extensively used by modelers, might not be a real plausible mechanism. Instead, we could have the following case: many populations cooperate, one being the observer and the others its input (simplified case). The observer population could learn to interpret only the input that is coherent with its current context, disregarding the other inputs. We could say that the relevant input wins the observer's attention, although no explicit competition takes place between the input populations (like lateral inhibition, etc).

Populations of neurons can cooperate and influence each-other’s activity, having either a driving effect (if one population has a high impact on another one) or just a modulator effect (slightly changing the dynamics of the target population, without changing its overall response). In order to properly define coherence, we need to make a clear distinction between simple association / combination of activities and coherent cooperation. While in the case of simple association the target population is always influenced by the input to produce a given response, in the case of coherent activation the target population will only respond if its interactions with the context populations “allow it”. To be more precise, for the same input stimulation, the target population could either respond or not, depending on the activity of the context. When the target responds to the input given the context, then we say we have a “coherent interpretation of the input in the present context”.

The principle of coherence is just a natural way of describing the dynamical activity of a complex system like the brain, composed of many subsystems that cooperate and interact. While it does not introduce radically novel concepts, it provides the framework for a new way of approaching brain / large-scale model research. The coherence principle predicts that a target subsystem (like a micro-column or microcircuit) couples to other subsystems in a non-trivial way, and that this coupling is mainly dependent on the dynamical state of the subsystem and of its associated context subsystems. For such a system to be flexible and generic enough, it seems plausible that the target subsystem has to be able to engage in different “trajectories” in time (by trajectory we mean dynamical evolution in time), depending on the interactions with the context. Taking the same subsystem, it would be of great importance that one and the same neural architecture should be able to exhibit a very large set of dynamical behaviors, as independent from the underlying architecture as possible with no trivial attractors or preferred states. It has been suggested that such systems could be formalized in terms of “far-from equilibrium computing devices” [14]. Put in a different way, the same neural subsystem would be able to “run different neural programs” at different times. In the case of the brain, the canonical micro-columns also called microcircuits could represent such subsystems.

From a different perspective, if information could be encoded in the trajectory in time of the subsystems, then it seems also plausible that the subsystems should be able to get engaged on corresponding trajectories by incomplete, partial information. This is essential if the system has to deal with an inaccessible environment (as is the case with the brains and the natural environment). Then, the subsystems cannot be “at the mercy of the input” drive from the stimuli, but must be able to predict, i.e. be independent from the input drive and reproduce as well as possible the trajectory associated with the partial stimulation. Such subsystems must then exhibit self-sustained activity, with trajectories merely initiated/constrained by the input. We might think of this in terms of predictive coding [15] but this time from a dynamical perspective.

The same stimulation of a subsystem at different times can lead to completely different responses, depending on the recent history and the instantaneous contexts influencing the subsystem. However, the most important aspect here is that such a subsystem can be rendered non-responsive to an input (non-coherent) if the context interactions bring it to a “dynamical refractory” state. If the subsystem is self-sustained (doesn’t need external energy to be active) and in addition it has no preferred states / attractors, then in principle it should be possible for a “context” to “modulate” the subsystem’s trajectory such that an observer population can recover the dynamical properties of the “neural program” running on the microcircuit at any given point in time. In this respect, the microcircuit is meta-stable in terms of dynamics, depending on the associated context modulation.

In this paper, we investigate a scenario of a self-sustained microcircuit stimulated by an input and “modulated” by a context. We show that an output population of neurons having synapses with spike-timing dependent plasticity can be used in principle to recover context sensitive trajectories of the microcircuit.

3. A Microcircuit "at work"

Studying neural microcircuits with rich dynamics, close to their biological counterparts, requires computationally effective models. One such neural model, developed by Izhikevich [16], has a very efficient implementation and it is, at the same time, very rich in dynamics. At the synaptic level, Senn and Markram have developed a biologically plausible algorithm for modifying the neurotransmitter release probability in a spike-timing dependent way [17].

Based on such plausible building-blocks, we have constructed a simulated neural architecture that consists of 4 distinct populations of neurons (Figure 2).

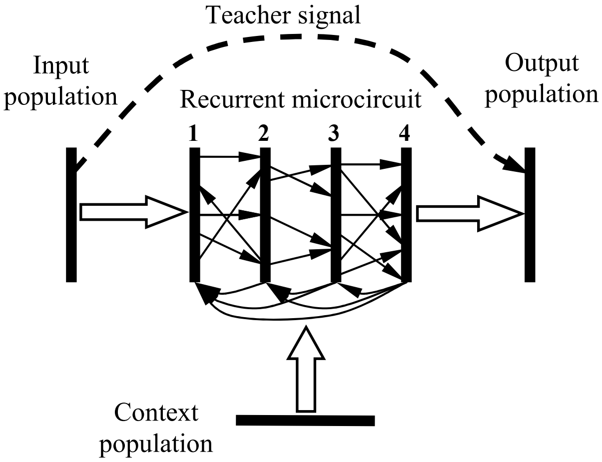

Figure 2. The architecture of the test model. An input population of 10 neurons randomly projects to layer 3 of a 4-layered neural microcircuit (10 neurons in each layer). An output and a context population, both consisting of 10 neurons, are connected with all neurons in the microcircuit (all layers). The context population sends axons to the microcircuit, while the output population receives input from all the neurons of the microcircuit.

Neurons. We used 3 different types of neurons, distributed differently, depending on the population (see "Architecture" below). The neural dynamics are modeled with a system of two coupled differential equations as described by Izhikevich, having parameters (*a, b, c, d*) that allow for the modeling of neurons with different properties (such as "regular spiking", "resonator", "chattering", etc) [16]. We have used "regular spiking" neurons for the input, output and context population neurons (see ref. [16], parameters: *a = 0.02, b =0.2, c = -70, d=8*). For the other population (the recurrent microcircuit), two additional neuron types are used: "resonator" and "chattering" neurons (see ref. [16], parameters for "resonator" neurons: *a = 0.1, b =0.26, c = -70, d=8;* parameters for "chattering" neurons: *a = 0.02, b =0.2, c = -50, d=2*).

Synapses. The synapses are modeled by specifically defining the dynamics of a synaptic conductance, g_{syn}. They are dynamic in the sense that the postsynaptic current (*PSC*) is delivered using the following set of equations:

$$PSC(t) = A_{syn} \cdot W_{syn} \cdot g_{syn}(t-1) \cdot (E_{syn} - U_{post}(t-1)) \tag{1}$$

$$g_{syn}(t) = g_{syn}(t-1) \cdot e^{-1/\tau_{syn}} \tag{2}$$

$$if\,(presynaptic\,spike)\,then\;\; g_{syn}(t) = g_{syn}(t-1) + 1 \tag{3}$$

where, A_{syn} is the absolute maximal synaptic strength, W_{syn} is the synaptic efficacy (between 0..1), g_{syn} is the time dependent synaptic conductance, E_{syn} is the reversal potential for the synapse type (-90 mV for inhibitory and 0 mV for excitatory synapses), U_{post} is the membrane potential of the postsynaptic neuron. Each time a presynaptic spike occurs, the value of g_{syn} is incremented by 1, allowing for the superposition of presynaptic stimulation effects.

Using this basic synaptic model, it is possible to represent both "fixed" and "plastic" synapses. In the case of fixed synapses, the value of the synaptic efficacy W_{syn} (which represents the presynaptic release probability) is not changing during the simulation. These synapses are instantiated with a fixed value for W_{syn} when the simulation is started. For the other case, of plastic synapses, the synaptic efficacy changes dynamically during the simulation, according to some rule. Synapses featuring spike-timing dependent plasticity (STDP) have an underlying algorithm for changing W_{syn} as a function of the pre- and postsynaptic spike timing. In the present study, we used the algorithm described by Senn and Markram [17].

Architecture. The first population, also called *input*, consists of a single sheet of 10 neurons that receive external current drive, injected by the experimenter. The neurons are modeled as “regular spiking” neurons (see ref. [16], parameters: $a = 0.02, b = 0.2, c = -70, d=8$). Each neuron sends projecting axons to neurons in the next population, namely to the third layer of the neural microcircuit (see Figure 2). Unlike the case of Dale’s law, the neurons modeled here are allowed to make both inhibitory and excitatory connections with postsynaptic targets. On average, each input neuron has roughly 5 synaptic contacts to postsynaptic neurons in the microcircuit. The projections are random, in the sense that the postsynaptic target is chosen randomly with a uniform probability of 5%. From these synapses, 20% are inhibitory and 80% excitatory. Their parameters are as follows: $A_{syn} = 0.05$, $\tau_{syn} = 15$ ms, $E_{syn} = -90$ mV for inhibitory synapses, and $A_{syn} = 0.02$, $\tau_{syn} = 25$ ms, $E_{syn} = 0$ mV for excitatory synapses. The synaptic efficacies are fixed (synapses are not plastic) to a value $W_{syn} = 0.7$.

The input population projects onto the third layer of a four-layered *recurrent microcircuit* with random connectivity (Figure 2). Each layer contains 10 neurons, with of a mix of 80% resonating ($a = 0.1, b = 0.26, c = -70, d=8$) and 20% “chattering” ($a = 0.02, b = 0.2, c = -50, d=2$) neurons [16], randomly distributed across layers (uniform distribution). This combination renders the microcircuit self-sustained even in the absence of external input. Self-sustained activity is possible because of the resonating properties of neurons (80% are of resonating type) which have low frequencies of resonance in the range of 5-30 Hz. The energy stored in the activity of the system gets amplified by the excitatory synapses, while the system avoids saturation because of the reluctance of neurons to high firing frequencies.

Within the microcircuit, the synapses randomly project between neurons but still, retinotopy is maintained to some degree. A synapse from a neuron at a given location usually projects to a neuron in a different layer but in a local neighborhood. Projections between layers are feed-forward (from low index to high index: eg. 1-2, 1-3, etc) as well as feed-back (from high index to low index: eg. 4-3, 4-2, 3-2, etc). Also, lateral interconnections between neurons in the same layer exist. Connectivity is mainly inspired from the rules of connectivity determined for the visual system [18]. The synapses are fixed and have the same parameters as described for input-to-microcircuit ones, except the fact that the synaptic efficacies are random: $W_{syn} \in (0..1)$.

The *output population,* consisting of 10 “regular spiking” neurons, is connected through synapses featuring STDP to every neuron in the recurrent microcircuit. We used for STDP the algorithm of Senn and Markram [17] which modifies the synaptic efficacy W_{syn} as a function of the relative timing between pre- and postsynaptic spikes. If the presynaptic spike occurs before the postsynaptic spike, in a given time window (50 ms), the synapse is potentiated, otherwise it is depressed. The amount of potentiation and depression depends on the difference between pre- and postsynaptic spiking times [17]. STDP synapses in our model have the following parameters: $A_{syn} = 0.05$, $\tau_{syn} = 20$ ms, $E_{syn} = 0$ mV. The synaptic efficacy W_{syn} is modified according to the aforementioned STDP algorithm and it is low bounded to a limit value of 0.05 (the release probability cannot drop to 0).

Also, there are a number of synapses with the input population that are activated only during the so-called "self-training" (when an external stimulus is applied to the input). Each neuron in the output population is connected with exactly one neuron in the input population, in an ordered fashion. Synapses are fixed, with parameters: $A_{syn} = 0.05$, $\tau_{syn} = 25$ ms, $E_{syn} = 0$ mV, $W_{syn} = 0.5$. Hence, the input is injected also into the output population neurons such as to provide a teacher signal. Such a signal is necessary since we want the output population to associate the activity of the microcircuit with the pattern that is presented as input during stimulation. In other words, the output population has to learn (through self-organization) to produce consistent activity patterns in response to the complex, self-sustained dynamics of the microcircuit. These activity patterns are constrained in our experiment to resemble the patterns provided as input during stimulation. However, this is only a simplifying design that aims at reducing the complexity of output patterns. In a more general case, the teacher signal might come from other microcircuits that have their own dynamics. For the sake of simplicity however, we provide an explicit teacher signal (the input in this case). Alternatively, the output population might as well be part of another dynamic microcircuit or the same microcircuit (this is a clearly problematic separation).

An additional population, that in a more general framework should be another microcircuit, is provided in order to inject a wave of spikes into the microcircuit. This population is called "*the context population*" and consists of 10 “regular spiking” neurons. The context population is controlled by injecting external currents during the experimental procedure. Neurons in the context population are connected via weak inhibitory synapses to all the neurons in the microcircuit population, with fixed synaptic efficacies, $W_{syn} = 1$. Their parameters are: $A_{syn} = 0.02$, $\tau_{syn} = 25$ ms, $E_{syn} = -90$ mV.

3.1. The dynamics of the model

The membrane oscillations of the “resonator” neurons, the “chattering” neurons and the highly recurrent synaptic connections, endow the microcircuit with a self-sustained activity, even in the absence of input (Figure

3). We call the activity of the microcircuit at any given time, the microcircuit's "mental state". Any spikes from the input or context population generate a perturbation of the microcircuit's "mental state".

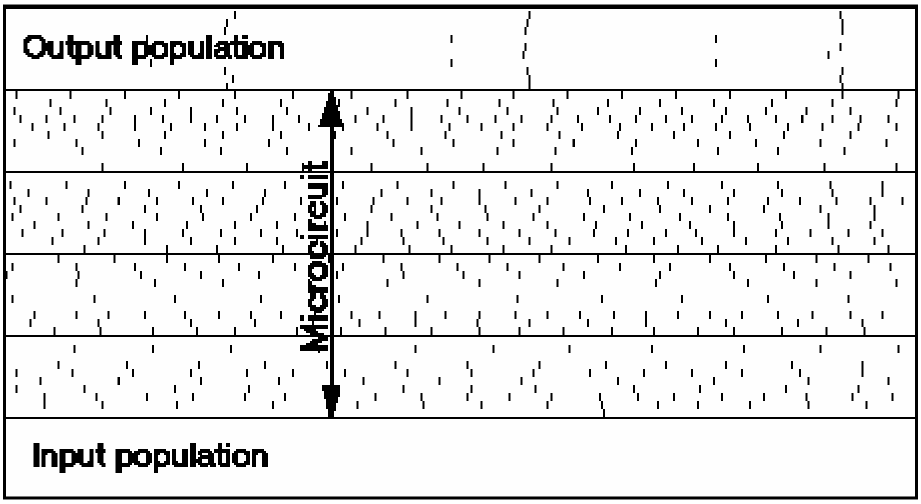

Figure 3. Self-sustained, spontaneous activity of the microcircuit in the absence of input. Time is represented on the "x" axis while each band on the "y" axis represents one layer of 10 neurons from the populations indicated on the plot. Each dot represents a spike of the corresponding neuron. The length of the trace is 500 ms.

An important observation is that the processing performed by the microcircuit is continuous. Between two input stimulations, the microcircuit has a sustained, continuous "mental activity". Unlike traditional neural network models, where the input or the background currents drive the activity (which does not exist in the absence of the input), the use of recurrent microcircuits opens a new, more plausible approach to neural processing. There is no external energy injected into the system while it is self-sustained.

Surprisingly, the spike trains generated by the recurrent microcircuit are highly irregular, resembling a spontaneous random activity (although, no random currents are injected, like in other models).

3.2. Experimental setup

Even for this relatively small system, with only 60 neurons, a clear and robust analysis method is hard to define. Using rate analysis on the neurons inside the microcircuit is impractical and should be avoided since we are interested in the millisecond dynamics and each individual spiking. An empirical observation during experiments was that although the rate of a neuron did not significantly change, for very weak inputs, the activity patterns of the output population dramatically change.

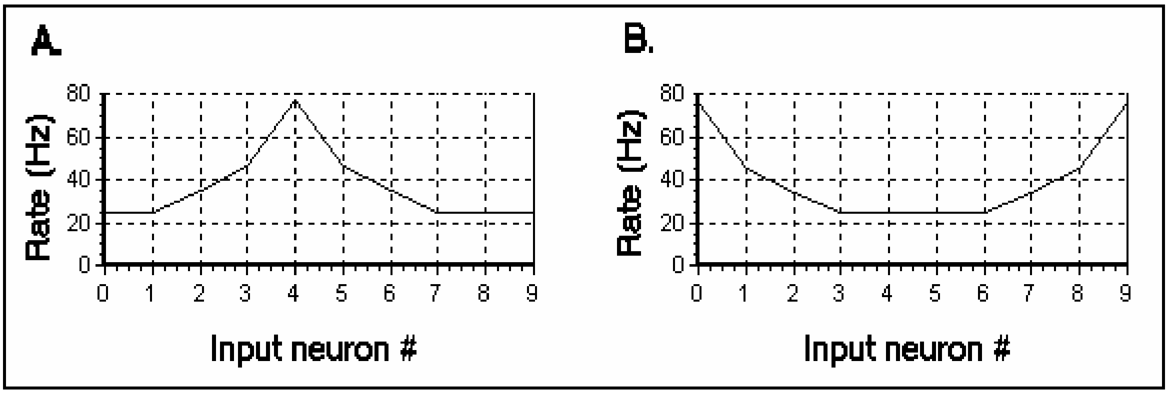

Figure 4. The two types of input signals used in the experiments. A. "Convex" rate profile; B. "Concave" rate profile. Each neuron fires a regular spike train with a firing rate corresponding to the value indicated in the plot.

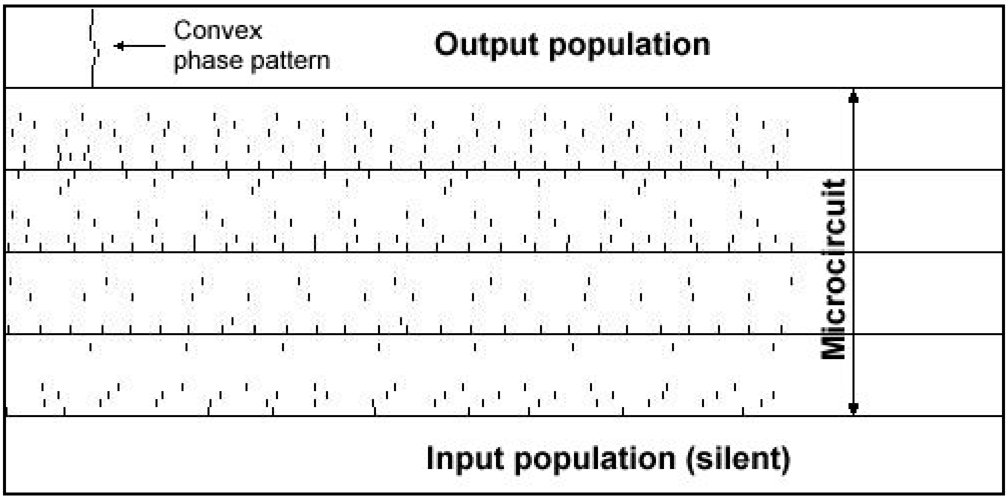

Figure 5. Example of "convex" phase coded pattern in the output population (top), after shutting down the inputs. The plot has the same elements as in Figure 3, but the activity is recorded at a different moment in time.

During experiments, a frequency-modulated signal is injected into the input layer. We used two types of signals: a "convex" signal and a "concave" signal (see Figure. 4), encoded in firing rates.

After presenting each input to the system, the activity pattern in the output population is analyzed. We empirically observed a consistent (and probably expected) transformation of the input rates into output phase coded responses. Consistent "convex" and "concave" activity profiles, encoded in the relative firing phase of neurons, emerge in the output population as a result of the "convex" and "concave" input rates, respectively. The phase coded response can be observed in Figure 5.

Rate analysis reveals no frequency-coded responses in the output population, not even during the stimulation period. The phase-coded responses appear only after shutting down the input. This suggests that the output population learns (through its STDP synapses) to interpret the microcircuit's spontaneous activity in a certain way ("convex" or "concave"). In other words the dynamical evolution of the microcircuit is mapped onto consistent patterns of output activity, coded as phase-patterns.

As a measure of "convexity" of the output responses, we used the following formula:

1. in 500 ms bins, a "convexity" measure (CM) is computed (for each "convex" event 20 was added to the "convexity" measure, for each "concave" event, 20 was substracted from CM);
2. an event was only validated if all the output neurons fired in a window of 20 ms length; if the relative phases of the firings respected an approximated "convex" profile, the event was considered to be "convex"; similarly, if the profile was approximated as "concave", a "concave" event was recorded.

During the first experiment, a "convex" input has been presented for about 40 seconds. The same experiment has been repeated also with a "concave" input. After presenting the rate patterns, the input is shut down and the output population is recorded (during this time the system receives no external stimulation).

Results, presented in Figure 6, clearly show that the output population interprets the reverberating activity (after stimulus shutdown) in a consistent way, suggesting that the "convexity" measure is relevant in analysis. Moreover, the high irregularity of recall events emphasizes the complexity of the neural activity and the richness of its dynamics (we should bear in mind that no random background currents are included into the model).

More intriguing results are obtained if the system is presented with both "convex" and "concave" inputs in two successive training epochs during the same experiment. The behavior of the system is in this case highly unpredictable. Alternating, episodic recalls of "convex" or "concave" events occur with high irregularity (or it might be that the analysis method fails to reveal some regularity). If the second training epoch is prolonged, the system's recall will become deterministic, favoring the last stimulus. We should mention that in the absence of any past input (training epochs), the system behaves randomly but eventually settles to a 0 "convexity" recall after a few seconds of activity (no preference for any output patterns).

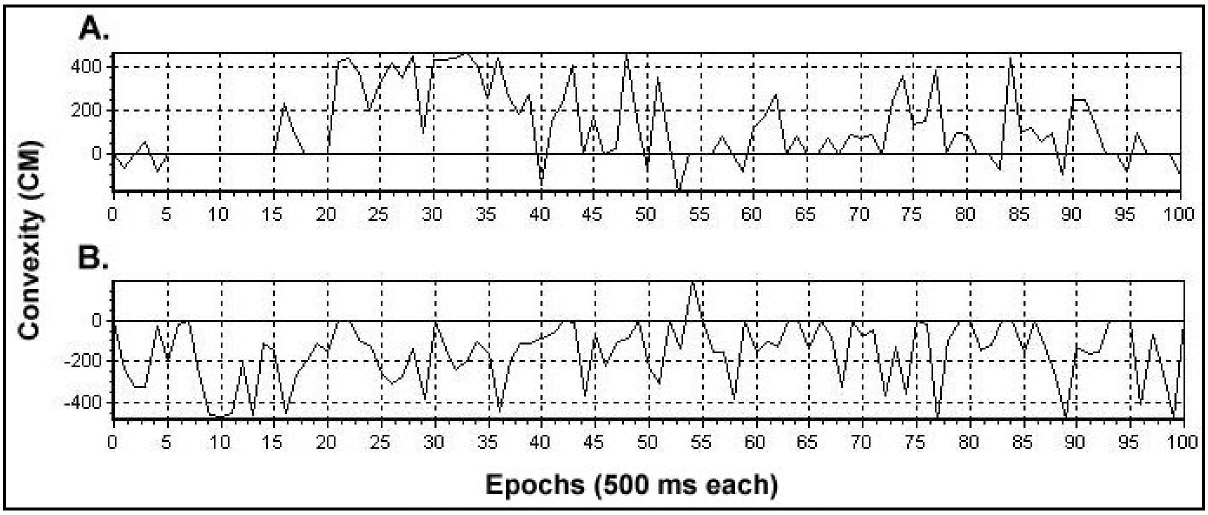

Figure 6. **"**Convexity" measure for the "convex" task (A) and "concave" task (B). Recall events are highly irregular in time, but consistent (always, with few exceptions due to the measuring method, convex or concave).

Finally, the most important experiment tries to reveal some possible mechanisms of context modulation and conditioning. First, the context population is activated (firing oscillating random spike trains in the gamma band, with a mean frequency of 30-50 Hz) during the "convex" training epoch. Then, it is shut down. After this step, the system is allowed to freely evolve for 50 seconds. A "concave" training epoch follows for 60 seconds. After the two training epochs, the inputs are shut down and the activity of the output population is observed.

In the beginning, the recall is mainly biased towards the last type of stimulus presented during the training epochs ("concave"). When the context population is activated, the associated stimulus is "recalled" (the "convex stimulus"). It seems that the STDP synapses are able to map two different trajectories of the microcircuit onto two different, even complimentary, firing patterns of the output. The microcircuit's trajectory is influenced by the "context" oscillations and this small modulation, although weak, can be recovered by the output neurons (Figure 7).

The implications of these findings are important since they show that a simple population of neurons is able to store and preferentially retrieve dynamic information presented in the past by using an association context. Synapses with biologically plausible models can self-organize to map different trajectories of a large system onto reliable consistent dynamics of another subsystem / population of neurons.

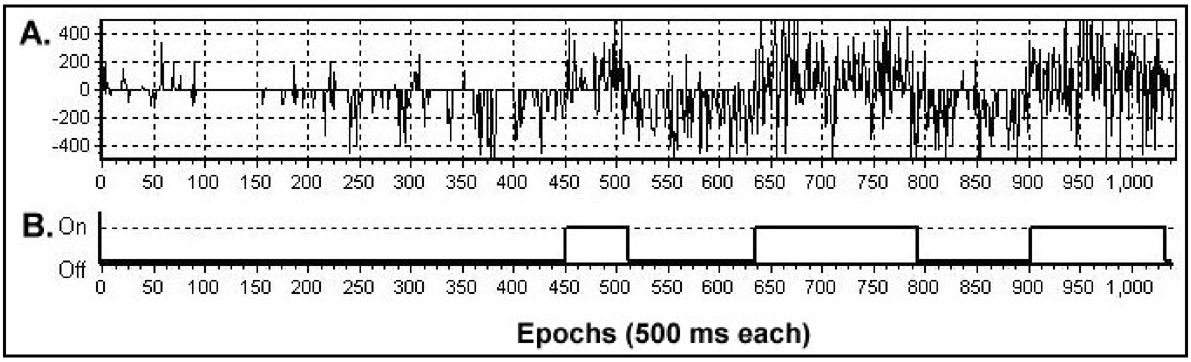

Figure 7. The system's behavior during the context experiment. A. The "convexity" measure of the output activity. B. The activation state of the "context population". Conditioned recall of "convex" patterns can be observed when the context population is "on". The plot monitors the "convexity" measure immediately after the end of training.

4. Conclusions

We have shown that even for the case of a relatively small neural microcircuit, the activity patterns that are produced are highly complex and dynamic. First of all, in our models we tried to produce a more plausible framework in which to study such systems. Therefore, we created a self-sustained microcircuit with ongoing dynamics and estimated how its trajectories can be characterized by a small set of readout neurons that self-organize their synaptic weights in order to capture some relevant associations. It is shown here that self-sustained activity can be mediated by the resonance properties of neurons, which allow for the stabilizing of the system. Because of the persistent activity of the circuit, the input stimulation represents a perturbation rather then a driving force. Moreover, the patterns of activity produced in the microcircuit, even in the absence of any external stimulation and any background noise, are irregular. This suggests that the highly complex dynamics of real neural systems should not be simply assumed to be noisy, but we should assess the degree to which they are subject to real noise.

Concerning the readout population analysis, some aspects need to be pointed out. Unlike other studies [7], we do not assume a readout that is explicitly trained, but rather use the information extracted collectively by a reduced set of neurons (phase patterns of 10 neurons in our case). We suggest that the activity of the 10 output neurons is a reduced characterization of the microcircuit's dynamics. Moreover, it is possible that the phase patterns produced by these neurons are able to recover fine temporal structure in the activity of the microcircuit, since we have not observed any robust rate changes in the circuit during the self-sustained, ongoing activity. Empirical observations led us to the conclusion that the resonating neurons studied here are preventing rate modulations because of their reluctance to high firing frequencies.

Finally, our results suggest that STDP is a versatile mechanism that can contribute to complex dynamical association processes. The plasticity rule, not only produces LTP and LTD, but it rather contributes to a dynamical extraction of relevant information from high dimensional transient states. This extraction can be modulated by subtle changes in the global state of a microcircuit, in our case induced by the activation of context inhibitory oscillations. We conclude that such interactions might form the basis for coherent and cooperative dynamics of large neural systems that finally give rise to memory, conditioning and maybe behavior.

Acknowledgements: We would like to thank to SC. NIVIS SRL (www.nivis.com) for supporting part of this research, and to Prof. Leonardo Chelazzi from the University of Verona. We would also like to thank Prof. Eugene Izhikevich from the Neurosciences Institute, San Diego, for useful comments on the manuscript and for his support.

References

[1] B.J. Scholl, A.M. Leslie, "Minds, Modules, and Meta-Analysis", *Child Development*, Vol. 72, 3:696-701, 2001.

[2] C.F. Stevens, "How cortical interconnectedness varies with network size", *Neural Computation*, 1:473-479, 1989.

[3] N. Brunel, V. Hakim, "Fast global oscillations in networks of integrate and-fire neurons with low firing rates", *Neural Computation*, 11:1621-1671, 1999.

[4] S. Fusi, P. Del Giudice, D.J. Amit, "Neurophysiology of a VLSI spiking neural network: Lann2", *In Proceedings of IJCNN 2000,* 2000.

[5] W. Gerstner, W.M. Kistler, *Spiking neuron models*, Cambridge University Press, Cambridge, UK, 2002.

[6] W. Maass, T. Natschläger, H. Markram, "Computational models for generic cortical microcircuits", In J. Feng, editor, Computational Neuroscience: A Comprehensive Approach, chapter 18. Chapman & Hall/CRC, Boca Raton, 2004, pp. 575-605.

[7] W. Maass, T. Natschläger, H. Markram, "Real-time computing without stable states: A new framework for neural computation based on perturbations", *Neural Computation*, 14(11):2531-2560, 2002.

[8] T. Natschläger, W. Maass, "Spiking neurons and the induction of finite state machines", *Theoretical Computer Science: Special Issue on Natural Computing*, 287:251-265, 2002.

[9] H. Jaeger, "The echo state approach to analysing and training recurrent neural networks", *GMD Report 148, German National Research Center for Information Technology*, 2001.

[10] L. Steels, R. Brooks, *The artificial life route to artificial intelligence: Building embodied, situated agents.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1995.

[11] R. Pfeifer, C. Scheier, *Understanding intelligence*. MIT Press, Cambridge, MA, 1999.

[12] R.V. Florian, Autonomous artificial intelligent agents. *Technical Report Coneural-03-01, http://www.coneural.org/reports/Coneural-03-01.php,* 2003.

[13] R.C. Mureşan, "The Coherence Theory: Simple Attentional Modulation Effects", *Neurocomputing*, Vol. 58-60C, Special Issue: Computational Neuroscience: Trends in Research 2004 Edited by E. De Schutter, pp. 949-955, 2004.

[14] R.V. Florian and D. Dumitrescu, Far-from-equilibrium computation. *Technical Report Coneural-04-01*, http://www.coneural.org/reports/Coneural-04-01.php, 2004.

[15] R.P.N. Rao, D.H. Ballard, Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects, *Nat. Neurosci.* 2 (1999), 79-87.

[16] E.M. Izhikevich, "Simple model of spiking neurons", *IEEE Transactions on Neural Networks*, 14:1569-1572, 2003.

[17] W. Senn, H. Markram, M. Tsodyks, "An algorithm for modifying neurotransmitter release probability based on pre-and postsynaptic spike timing", *Neural Computation*, 13:35–67.

[18] Patricia S. Churchland, T.J. Sejnowski, *The computational brain*. MIT Press, Cambridge, MA, 1999.

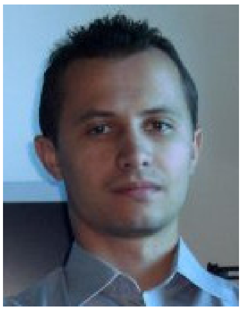

Raul C. Mureşan is a researcher at Frankfurt Institute for Advanced Studies. His research interests include modeling of large scale neural systems, the study of biologically plausible neural microcircuits and bio-inspired artificial intelligence. (Home page: http://www.raulmuresan.home.ro)