Associative memory using action potential timing

Roland Müller, David J.C. MacKay† and Andreas V.M. Herz*

*Institut für Theoretische Physik, Universität Bremen, 28334 Bremen, Germany;
and Innovationskolleg Theoretische Biologie, Humboldt-Universität zu Berlin,
Invalidenstrasse 43, 10115 Berlin, Germany [1]*
†*Cavendish Laboratory, Cambridge, CB3 0HE, United Kingdom*

Abstract

The dynamics and collective properties of feedback networks with spiking neurons are investigated. Special emphasis is given to the potential computational role of subthreshold oscillations. It is shown that model systems with integrate-and-fire neurons can function as associative memories on two distinct levels. On the first level, binary patterns are represented by the spike activity — "to fire or not to fire." On the second level, analog patterns are encoded in the relative firing times between individual spikes or between spikes and an underlying subthreshold oscillation. Both coding schemes may coexist within the same network. The results suggest that cortical neurons may perform a broad spectrum of associative computations far beyond the scope of the traditional firing-rate picture.

Introduction

A significant fraction of the communication between single neurons is based on action potentials, short pulses of electrochemical activity. The biological relevance of the exact temporal organization of action potentials has been demonstrated in many systems, including the sonar system of bats [1], the auditory system of owls [2], the olfactory system of locusts [4], and the visual system of cats [5, 6] and monkeys [7].

Most traditional neural network models, in contrast, neglect the exact timing of action potentials and describe the output of a cell as a continuous variable which is usually interpreted as a "firing rate" — the time-averaged rate of action potentials generated. This output variable is slowly varying in time relative to the interval between action potentials of a single cell. There are instances where such a simplified approach might be justified on a phenomenological network level [8, 9]. However, the above experimental results strongly suggest models where computations are based on the timing of individual action potentials.

There are also various theoretical reasons for studying such models. First, consider the information content of a train of action potentials. Even complex cognitive tasks are often accomplished within a few hundred milliseconds [10] although they involve multiple cortical regions typically operating at firing frequencies below 50 Hertz — which means that each subnetwork has to complete its computation within a small number of firing cycles. However, it will take about 100 cycles to convey a firing rate accurate to 10% if we assume a firing-rate code with Poisson-distributed interspike intervals; in contrast, the interval between two spikes (of the same or different neurons), or between a spike and a background oscillation, is a real number that can be identified, using neurobiological mechanisms, to a

---

[1]Permanent address

precision better than a millisecond. Such real numbers could be conveyed once every cycle of a background oscillation, that is, every 20 milliseconds for gamma-oscillations and every 100 − 150 milliseconds for theta-oscillations. Certainly enough, distributed population codes may greatly improve the accuracy of a firing-rate scheme, but they improve the precision of a spike-time code at least equally well. It thus seems plausible that information processing schemes that use spike timings to represent continuous quantities could be many times more efficient than traditional schemes based on firing rates.

Second, a system using a spike-time code might allow useful computations to be performed more naturally than a rate-coding network. Hopfield [11] has proposed a model of pattern recognition in which stimulus intensities are recoded into time advances of spikes relative to a background oscillation. This is done as follows: Consider a neuron subject to a subthreshold oscillation. If, in addition, the neuron receives a constant excitatory input current the subthreshold oscillation will be moved towards more depolarized values of the membrane potential. For large enough current the oscillating membrane potential crosses the firing threshold and the neuron will fire. The larger the input current — i.e. the stronger the stimulus — the earlier the neuron fires. This means that an analog pattern of stimulus intensities can be represented by an analog pattern of relative spike times.

Such a spike pattern could be recognized by a downstream coincidence detector with a short integration time constant and appropriate firing threshold. This cell will fire if the action potentials of the spike-encoding neurons arrive simultaneously. It will therefore be capable of recognizing the pattern if appropriate time delays between the spike-encoding cells and the coincidence detector compensate the relative timing of the action potentials. A memorized pattern is thus stored in the time delays — not in the synaptic weights as in traditional networks — and read out by a coincidence detector which functions as a "grandmother neuron" for the specific pattern. The synaptic weights do not have to be precise any more. A whole set of patterns can be processed if there exist multiple grandmother cells, each dedicated to a particular pattern.

The entire scheme has some advantages which are difficult to achieve in a network with rate coding. For example, in the special case where the time advance of a spike is given by the logarithm of the stimulus intensity, the pattern of relative firing times will be unchanged if the intensities of all stimuli are increased by a constant factor — the entire firing pattern will simply occur earlier relative to the background oscillation. So, on the one hand, the recognition of the pattern will be intensity invariant, while on the other hand, the information about the intensity will still be retained in the timing of the grandmother cell's response. Furthermore, all components of the input pattern are equally important for the recognition, independent of the absolute size of the components. Additionally, a pattern can be broken into smaller sub-parts to be processed in parallel. The scheme also allows a response to be built up hierarchically: a super-pattern (a combination of patterns in particular ratios) can be detected by a super-grandmother cell which receives inputs from the corresponding grandmother cells with appropriate time delays.

Third, the use of time to represent continuous quantities liberates the firing rate to serve a new purpose. If for example, in olfaction, the spike timing of a grandmother cell represents the intensity of an odor, the firing rate might represent the probability that the odor is really there.

Overview of this article: Although Hopfield's scheme offers a number of fascinating computational capabilities, it is based on a purely feedforward network architecture where

quasi-stationary stimuli are processed by a first layer of spike-encoding cells and a second, downstream layer of grandmother cells which do the read-out. In the present article, this scheme is extended to feedback networks where the spiking neurons also provide recurrent input to each other. In the next section we first describe the dynamics of integrate-and-fire neurons and traditional firing-rate models. In the two subsequent sections we then show how networks of integrate-and-fire neurons can function as autoassociative memories for binary and analog patterns, respectively. In the first case, a binary spike-activity code results on the network level, in the second the exact timing of action potentials is used to process time-encoded analog information. We close our considerations with a discussion of the results and some further ideas.

Neuron models

Various brain regions have been hypothesized to function as autoassociative memories, for example the CA3 region in the hippocampus, or locally connected systems of pyramidal cells in neocortical association areas [12, 13, 14, 15]. Traditionally, these systems have been modeled using neurons whose output is described by a short-time averaged firing rate [16, 17, 18, 19, 20, 21].

Autoassociative memories are capable of storing a set of patterns in the sense that they complete or correct a pattern when presented with an incomplete or noisy input. To do so, the networks store correlations within the patterns in the synaptic weights. The weights are determined by learning rules such as the Hebb rule. By this procedure, learned patterns are implemented as fixed-point attractors of the network dynamics. Starting from an incomplete or noisy initial state, the system relaxes to a nearby fixed point — the retrieval of the memory.

Extensive theoretical results have been derived regarding the convergence properties and storage capacity of such models. However, by definition there are no action potentials in firing-rate models and thus there is no possibility for temporal codes in the millisecond range. From the arguments given in the Introduction, the biological relevance of these models therefore appears to be questionable, at least as far as rapid information processing is concerned. It is thus important to investigate biologically more realistic networks with spiking neurons which may support a temporal code based on the firing of individual action potentials [11, 22, 23, 24, 25, 26].

Biological neurons fire an action potential when their membrane potential u reaches a threshold value $u_{threshold}$. The action potential then propagates along the axon to synapses on the dendritic trees of postsynaptic neurons. When the action potential arrives at a synapse after some axonal delay time it initiates the release of neurotransmitters which lead to a flow of ionic currents that depolarize (excitatory synapse) or hyperpolarize (inhibitory synapse) the postsynaptic cell. Depending on the integrated inputs from many thousands presynaptic cells a neuron is typically connected with, the postsynaptic cell will in turn fire an action potential at some time and thus influence further neurons.

Formal neural networks are constructed to capture important features of these intricate dynamical processes. In order to allow for an analysis of the collective dynamics in large feedback networks, the microscopic dynamics have to be simplified. For example, individual neurons are often considered to be electrotonically compact and described by a single dynamical variable, the membrane potential at the soma of the neuron. Models differ in the

way how they represent both neural activity and the response to input from other neurons.

Integrate-and-fire neurons: Integrate-and-fire model neurons operate as leaky integrators as long as the membrane potential u_i is below the firing threshold $u_{threshold}$,

$$C\frac{d}{dt}u_i(t) = -\frac{1}{R}(u_i(t) - u_{rest}) + \sum_j T_{ij} \ PSC_j(t) + I_i^{ext}(t). \qquad (1)$$

Here C denotes the capacitance, R the resistance, PSC_j a unitary postsynaptic current which, when multiplied by the synaptic weight T_{ij}, describes the effect of the firing of neuron j on i's membrane potential. In equation (1) active dendritic processes are neglected so that the input can be approximately described by a linear summation of the individual synaptic terms, together with an input $I_i^{ext}(t)$ from outside the network. When the membrane potential of cell i reaches $u_{threshold}$, the cell instantaneously produces a uniform δ-shaped action potential and resets its membrane potential to u_{rest}. The output of an integrate-and-fire neuron is thus a sequence of δ-pulses.

Graded-response neurons: Model neurons whose membrane potential changes according to the differential equation

$$C\frac{d}{dt}u_i(t) = -\frac{1}{R}(u_i(t) - u_{rest}) + \sum_j T_{ij} \ V_j(t) + I_i^{ext}(t). \qquad (2)$$

are called "graded-response neurons." The output of a neuron j is now represented by a short-time averaged firing rate V_j. To specify the dynamics one has to define an activation function that describes the graded response of the firing rate to changes of the membrane potential. The neural response to long, constant current pulses injected into the cell body has been recorded in vitro. The current-to-frequency curves describe the adapted firing rate as a function of the injected current and exhibit a sigmoid shape which is usually used for the activation function $V = g(u)$.

Neurons with discrete-time dynamics: Within an even more simplified approach, time advances in steps of fixed length Δt which may represent the total duration of a single action potential, roughly 1 millisecond. The output of a neuron is "updated" at these discrete times, and the evolution equations become discrete-time iterations,

$$V_i(t + \Delta t) = g(\sum_j T_{ij}V_j(t) + I_i^{ext}(t)). \qquad (3)$$

Such networks are also known as "iterated-map networks" [20]. In the extreme case of a threshold function which vanishes for negative u and is unity for positive u, the output of a neuron reduces to a two-state variable $S = 0/1$,

$$S_i(t + \Delta t) = \Theta(\sum_j T_{ij}S_j(t) + I_i^{ext}(t)). \qquad (4)$$

In the Little model [16], the updating rule (4) is done in parallel on the network level, in the Hopfield model [17], the neurons are updated sequentially, i.e. one-by-one.

Associative storage of binary patterns

In this section we show that under certain conditions networks with integrate-and-fire neurons are capable of storing binary patterns — on the attractor of the dynamics, single neurons either fire periodically (the "on"-state) or remain completely quiescent (the "off"-state). The result demonstrates that although the underlying time evolution strongly differs from that of graded-response neurons or binary neurons, on the level of firing rates, the collective network properties may again be quite similar. Our theoretical analysis also provides a connection to early simulation studies which revealed associative capabilities in networks with biologically more realistic spiking neurons [27].

Model with constant background activity: Consider a network of N integrate-and-fire neurons which are coupled via synaptic weights T_{ij}. An action potential of one cell arrives at a postsynaptic neuron after a fixed and uniform axonal time delay τ_a. In addition, let us assume that neurons receive a constant input current slightly less than $R(u_{threshold} - u_{rest})$ so that their membrane potential is raised to a value just below firing threshold. Within a biological context, this input can be thought of an approximation to randomly impinging inputs that put the cell in an "alert" state as it has been hypothesized for cortical neurons [28, 29]. Neurons will thus not fire without additional excitatory recurrent input from other cells within the network. The postsynaptic current is approximated by a δ-pulse which is justified if the time constants of the ion channels involved are very short.

At time t_0 the system is presented a binary pattern of input currents such that the membrane potentials of a group G_0 of neurons are raised above firing threshold whereas all other neurons receive no input. Subsequently the network exhibits the following dynamics.

After the firing of the initial spikes, the membrane potentials of the group G_0 are reset to u_{rest}. After the delay τ_a, the action potentials generated at time t_0 arrive at the postsynaptic neurons. By that time, neurons which fired at time t_0 have fully integrated the constant input current if the membrane time constant $\tau_{RC} = RC$ is much smaller than the time delay τ_a. Thus the membrane potentials of all neurons are again just below threshold. Upon arrival the action potential fired by neuron j at time t_0 triggers a postsynaptic current which instantaneously raises the membrane potential of neuron i by the amount T_{ij}. Neuron i will therefore fire at time $t_0 + \tau_a$ if $\sum_{j \in G_0} T_{ij} > 0$ and will stay quiescent if the sum is negative.

The above reasoning implies that at time $t_0 + k\tau_a$, where k is a positive integer, a certain group G_k of neurons fires. This group is the *same* group of neurons that is in the on-state ($S = 1$) in the k-th iteration of the Little model [16] started with the same initial conditions (same group G_0 of neurons with $S = 1$ at time t_0) and the same couplings T_{ij} which store a set of patterns. This implies that both networks retrieve the same pattern and that after typically 5 times the delay τ_a (5 iterations for the Little model), the group of neurons which fire (have $S = 1$) remains unchanged. There is only one though important difference between both scenarios. The Little model literally reaches a fixed point, i.e. the S_i remain unchanged, whereas in the present network the binary pattern is represented by a τ_a-periodic firing pattern. Note that a short membrane time constant τ_{RC} is necessary to guarantee that i) the constant input current is integrated after the time τ_a for a neuron which fired and reset to u_{rest} at time $t_0 + k\tau_a$, and ii) that the recurrent input $\sum_{j \in G_{k-1}} T_{ij}$ has decayed to zero after the time τ_a for a neuron which did not fire at time $t_0 + k\tau_a$. If one allows for larger τ_{RC} the dynamics will still be qualitatively similar to the dynamics of the

Little model. Independent of the value τ_{RC} the dynamics will be identical if one arranges several networks in a staggered loop structure. In such a loop the output of one network serves as input to the subsequent network and there are no recurrent connections within one network. The number of networks needed to close the loop (the last network in the loop projects back to the first network) depends on τ_a and τ_{RC}.

Model with subthreshold oscillation: So far, we have not taken into account subthreshold oscillations such as the hippocampal theta rhythm — the original motivation for this study. Let us now turn to such networks. If the input current $I(t)$ to an integrate-and-fire neuron contains an oscillatory component $I_{osc}(t) = I^* \cos(\omega t)$ one can simply solve (1) without this oscillatory input component and afterwards add an oscillatory membrane potential term $u_{osc}(t) = A \cos(\omega t + \phi)$ to the solution. Here $A = f(\tau_{RC}\omega)RI^*$, $\phi = g(\tau_{RC}\omega)$ and $\tau_{RC} = RC$. Without loss of generality, one may set $\phi = 0$ since ϕ can be absorbed in t_0. We therefore drop the term $I_{osc}(t)$ in the input current $I(t)$ and just add $u_{osc}(t)$ to the solution.

Let us again assume a constant input current $I(t) = I$. The membrane potential exhibits a subthreshold oscillation if I vanishes and A is not too large. An increasing input current I shifts the subthreshold oscillation to more depolarized values of the membrane potential. If I exceeds some critical value, the oscillating membrane potential will cross $u_{threshold}$ and the neuron will fire. The firing of neuron j triggers a rectangular postsynaptic current pulse with amplitude T_{ij} and duration τ_{PSC} which is added to the input current I_i of neuron i after a time delay τ_a.

Given a subthreshold oscillation with amplitude A and frequency $f = 1/P$ we want to store binary patterns in a network of N neurons which receive a constant input current I slightly less than $R(u_{threshold} - u_{rest} - A)$ so that the maximum of the membrane potential oscillation is just below $u_{threshold}$ as explained for the previous model. A binary pattern consists of the neurons i which fire ($S_i = 1$ in the Little model) or do not fire ($S_i = 0$) in the time interval τ_{fire} before the maximum of an oscillation cycle.

To store p uncorrelated patterns in the Little model using the Hebb rule one sets

$$\tilde{T}_{ij} = \frac{1}{N} \sum_{\mu=1}^{p} (2\xi_i^\mu - 1)(2\xi_j^u - 1). \tag{5}$$

The retrieval properties of the Little model will not change if the synaptic weights $\tilde{T}_{ij}$ are replaced by $T_{ij} = \kappa \tilde{T}_{ij}$, $0 < \kappa \le 1$.

The present system will operate properly as an associative memory if *i*) in a given cycle, neurons do not fire earlier than τ_{fire} before the maximum, and *ii*) τ_{fire} does not exceed $P/4$. For appropriate τ_{fire}, the first condition can be satisfied if κ is chosen small enough so that $\max_i(\sum_j |T_{ij}|) = \eta A$, where $\eta = \sin(\omega \tau_{fire})$. This is because the recurrent input raises the membrane potential by at most $\max_i(\sum_j |T_{ij}|) = \eta A$. We then choose

- $\tau_a = 2\tau_{fire}$ so that neurons do not receive any recurrent input in the time interval τ_{fire} after the maxima of the oscillation. Since the membrane potential decreases after the maxima, the neurons cannot fire in the interval τ_{fire} after the maxima.

- $\tau_{PSC} = P - \tau_{fire}$ so that all PSCs last at least until the maximum of the present oscillation cycle — the first PSC may be triggered as early as τ_{fire} after the maximum of the previous oscillation.

- $\tau_{RC} \leq -\frac{P-3\tau_{fire}}{\ln(\epsilon)}$ so that all PSCs are "integrated" until τ_{fire} before the maximum of the present oscillation cycle (the last PSC may not arrive till $\tau_a = 2\tau_{fire}$ after the maximum of the previous oscillation cycle). Here "integrated" is defined as $(1 - \exp(-\frac{P-3\tau_{fire}}{\tau_{RC}})) \geq 1 - \epsilon$ where ϵ denotes the relative accuracy to which the actual membrane potential change due to a single PSC is required to have relaxed towards the asymptotic value.

With the above choice of parameters the network will exhibit the following dynamics if initialized with a binary pattern of input currents such that the membrane potentials of a group G_0 of neurons are raised above $u_{threshold}$ at or slightly before the maximum of an oscillation cycle. After the firing of the initial spikes, the membrane potentials of group G_0 are reset to u_{rest}. After the delay τ_a, the action potentials arrive at the postsynaptic neurons and trigger rectangular current pulses PSC with duration τ_{PSC}. By the time τ_{fire} before the maximum of the next oscillation cycle all the PSCs will be integrated. A postsynaptic neuron i will fire if $\sum_{j \in G_0} T_{ij} > 0$ and will stay quiescent if the sum is negative.

In the time interval τ_{fire} before the maximum of the k-th oscillation cycle, where k is a positive integer, a certain group G_k of neurons fires. As in the case without subthreshold oscillation, this group is the *same* group of neurons that is in the on-state ($S = 1$) in the k-th iteration of the Little model [16] that was presented the same initial pattern and has the same couplings T_{ij}. This again implies that all results obtained for the Little model regarding fixed points, convergence, and storage capacity can be directly transfered to the present model.

The parameters τ_a, τ_{PSC} and τ_{RC} do not have to be identical for all neurons as long as the variation from the mean values defined above is not too large. The mean value for τ_a may be arbitrary if one arranges several networks in a staggered loop structure with appropriate phase shifts of the membrane potential oscillation.

Analyzing the correspondence between the present model and the Little model we have so far not considered any coding in the temporal domain — besides using the oscillation to generate a global timing mechanism. Binary-valued patterns are expressed in the firing or not firing of a neuron before the maximum of a given oscillation cycle. But potentially, there is additional information in the time at which a neuron fires relative to the maximum of the oscillation cycle. If a neuron i reaches $u_{threshold}$ it will fire at time $t_i^{fire} = \frac{1}{\omega} \arccos(\frac{1}{A} \sum_{j \in \mathcal{F}} T_{ij})$ before the maximum of the oscillation where $\mathcal{F}$ denotes the group of neurons that fired in the previous oscillation cycle. This means that the value of $\sum_{j \in \mathcal{F}} T_{ij}$ is encoded in the firing time of neuron i.

To recognize a pattern one could again use a coincidence detector as proposed by Hopfield [11] within a feedforward scheme. Any such grandmother neuron only receives input from the n neurons of the network which are active in the pattern the specific grandmother neuron is coding for. The cell fires when n action potentials arrive simultaneously. In the previous model where all active neurons fired simultaneously the cell would fire when the network has reached the fixed point the cell is coding for, as desired. However, the grandmother cell would also fire if the network has reached another fixed point in which not only those n neurons are active but also additional ones. This situation is very unlikely to happen in the network now described because a certain pattern is not just represented by its active neurons but also by the relative timing of their action potentials. Additional active neurons will now contribute to the postsynaptic potential of the n neurons and change their

relative firing times so that their action potentials do not any more arrive simultaneously at the grandmother cell which will therefore remain quiescent.

Associative storage of analog patterns

In the previous section a *binary* pattern was represented by the activity of all neurons just before the peak of the oscillation — a code of "to fire or not to fire." The exact time advance with respect to the oscillation peak was not used to represent the pattern. *Analog* patterns can be represented by time advances if one takes elementary biophysics into account. The size of a postsynaptic current PSC is *not* constant but depends on the present membrane potential — and thus in systems with subthreshold oscillations on the time relative to the underlying oscillation.

As an example, consider excitatory synapses in the central nervous system that use the neurotransmitter glutamate. Glutamate receptors can be broadly separated into two main types — those sensitive to NMDA and those that are not. If brief pulses of glutamate are applied under voltage-clamp conditions the current response of a non-NMDA synapse reverses near zero millivolt and exhibits a fairly linear relationship with respect to the membrane potential. This means that the size of a PSC that is mediated exclusively by non-NMDA receptors decreases with increasing membrane potential. The amplitude of a PSC triggered by an action potential which arrives on the upward slope of an oscillation cycle, is thus an increasing function of the time advance of that action potential.

For NMDA synapses, the situation is different for two reasons. First, in hyperpolarized cells, NMDA channels are blocked by Mg^{2+}-ions and do not pass much current. At more depolarized levels, the Mg^{2+}-blocks are removed and the size of the PSC increases — opposite to the case of non-NMDA synapses. This implies that on the upward slope of the oscillation the amplitude of a PSC is a decreasing function of the presynaptic time advance, but on the *downward* slope it is again an increasing function.

Second, NMDA channels have long opening times up to roughly hundred milliseconds. This means that the ionic current through a synapse will be an average over a significant fraction of the oscillation cycle and may thus not be used as a precise temporal marker. Nonlinearities may, however, lead to a better temporal precision than expected from this argument. Detailed computer simulations are needed to evaluate the potential of NMDA channels for temporal coding.

Inhibitory synapses have reversal potentials below the resting potential of a cell. As in the case of NMDA synapses — but with the advantage of short time constants — the size of a PSC is an increasing function of the time advance on the downward slope of the oscillation.

In the cases discussed above, the amplitude of the PSC is a monotone function f of the time advance x_j^k of neuron j in the k-th cycle. If we assume that individual contributions add linearly, the total excitation will be given by $a_i^{k+1} = \sum_j T_{ij} f(x_j^k + \tau_a)$. The postsynaptic neuron will fire the earlier, the stronger the cell is depolarized. Thus x_i^{k+1} is given by an increasing function g of a_i^{k+1} and we finally obtain

$$x_i^{k+1} = g\left(\sum_j T_{ij} f(x_j^k + \tau_a)\right). \tag{6}$$

Note that the x_i^{k+1} on the left-hand side of the equation refer to time advances in the next cycle.

The axonal time delay τ_a may vary for different neurons but has to be chosen appropriately so that the action potentials arrive on the right slope (upward or downward) of the oscillation. In case they have to arrive on the upward slope, τ_a has to be rather long at first sight. However, the scheme can be readily generalized to "staggered projections" as they are found in the hippocampus. In this case, the constraints on the axonal time delay are greatly reduced. If, for example, there is a 90-degree phase shift between two consecutive subnetworks (as between the dentate gyrus and the CA3 region), the output of one network can be directly fed into the next network with no or only very small delays.

Another possibility along this line — but within a single network — will arise if there are multiple subsets of neurons with different time delays. For simplicity let us consider the case of two subsets $\mathcal{A}$ and $\mathcal{B}$. In any given oscillation cycle, neurons in subset $\mathcal{A}$ fire before those in subset $\mathcal{B}$. Subset $\mathcal{A}$ receives input from the firings of subset $\mathcal{A}$ and $\mathcal{B}$ in the previous cycle, i.e. $\tau_a^{\mathcal{A}\to\mathcal{A}}$ and $\tau_a^{\mathcal{B}\to\mathcal{A}}$ is long. Subset $\mathcal{B}$ receives input due to the firing of $i)$ subset $\mathcal{B}$ in the previous cycle, i.e. $\tau_a^{\mathcal{B}\to\mathcal{B}}$ is also long, and $ii)$ subset $\mathcal{A}$ in the *present* cycle, i.e. $\tau_a^{\mathcal{A}\to\mathcal{B}}$ is *short*. This means that subset $\mathcal{A}$ performs a preliminary signal processing which is used for the computation done by subset $\mathcal{B}$ in the *same* cycle.

Equation (6) is formally almost identical to the evolution equation of an iterated-map network [20], the continuous-valued generalization of the Little model [16]. Indeed, in the most simple case with linear f (which may be justified for excitatory non-NMDA synapses and action potentials that arrive in the linear part of the upward slope), the dynamics reduce to the well-known iteration equations

$$x_i^{k+1} = g\left(\sum_i T_{ij}x_j^k + \theta_i\right) \tag{7}$$

which should be compared with equation (3). The general equation with nonlinear f will probably result in interesting new computational properties.

Discussion

The present study shows that model systems with integrate-and-fire neurons can operate as associative memories on (at least) two separate levels. On the first level, binary patterns are represented by the firing state of individual neurons — "on" or "off." This dichotomy is possible because in the systems studied, neurons approach an attractor where they either fire periodically or are completely quiescent. The relative firing times between individual spikes or between spikes and the underlying subthreshold membrane oscillation may be used to encode analog patterns, the second level of associative memory. Furthermore, both coding schemes can coexist in the same network.

The shape of a postsynaptic potential can also be used to do computations based on a spike-time code [26]. If for example, by the time the membrane potential of a postsynaptic neuron reaches firing threshold, all PSPs are still in their rising segment then the earlier a PSP was triggered the earlier the postsynaptic neuron will fire. Extending this idea to systems with subthreshold oscillations, an equation similar to equation (6) is obtained. However, it is an implicit equation for the time advance x_i^{k+1} because x_i^{k+1} also appears

inside the function f on the right-hand side. For linearly increasing PSPs the situation simplifies since in this case f is a linear function. If the system operates in the linear part of the upward slope, the function g will also be linear and the iteration equations become

$$x_i^{k+1} = \lambda \left(\sum_i T_{ij} x_j^k + \theta_i \right) \qquad (8)$$

where λ is some real number.

As shown in the last section, spike-time coding is facilitated in systems where the size of neural signals is *not* uniform but depends on the current state of the postsynaptic cell. In the example presented, the biophysics of synaptic transmission imply postsynaptic potentials whose amplitude directly reflects the present postsynaptic membrane potential. In any system where an oscillating component is added to the neuronal input current, the dependence on the postsynaptic membrane potential becomes a temporal relation, a potential basis for a spike-time code to process analog information.

In closing, let us emphasize that the voltage dependence of postsynaptic currents is just one example for this general coding principle. Any dependence of a parameter of the single-neuron dynamics on state variables of a cell may potentially serve a computational role within a temporal code. Other examples include the dependence of the membrane time "constant" τ_{RC} on the membrane potential [30] or the dependence of the synaptic weights T_{ij} on the recent firing history of a neuron [31].

References

[1] Suga, N., O'Neill, W.E., Kujirai, K., & Manabe, T. *J. Neurophysiol.* **49**, 1573-1626 (1983).

[2] Konishi, M. *The Harvey Lect.* **86**, 47-64 (1992).

[3] Carr, C.E., Heiligenberg, W. & Rose, G.J. *J. Neurosci.* **6**, 107-119 (1986).

[4] Laurent, G. & Davidowitz, H. *Science* **265**, 1872-1875 (1994).

[5] Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M. & Reitboeck, H.J., *Biol. Cybern.* **60**, 121-130 (1988).

[6] Gray, C.M. & Singer, W., *Proc. Natl. Acad. Sci. USA* **86**, 1698-1702 (1989).

[7] Kreiter, A. and Singer, W. Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *J. Neurosci.* **16**, 2381-2396 (1996).

[8] Amit, D.J. The Hebbian paradigm reintegrated — local reverberations as internal representations. *Behav. Brain Sciences* **18**, 617-626 (1995).

[9] Seung, H.S. How the brain keeps the eyes still. Preprint, AT&T Bell Laboratories, 1996

[10] Thorpe, S.J. Temporal synchrony and the speed of visual processing. *Behav. Brain Sciences* **16**, 473-474 (1993).

[11] Hopfield, J.J. Pattern recognition computation using action potential timing for stimulus representation. *Nature* **376**, 33-36 (1995).

[12] Marr, D. A theory for cerebral neocortex. *Proc. R. Soc. B* **176**, 161-234 (1970).

[13] Marr, D. Simple memory: a theory for archicortex. *Phil. Trans. R. Soc. B* **262**, 24-81 (1971).

[14] Rolls, E.T. *The Computing Neuron* (edited by R. Durbin, C. Miall and G. Mitchison) 125-159 (Addison-Wesley, Wokingham, 1989).

[15] Rolls, E.T. *Neurobiology of Comparative Cognition*, (edited by D.S. Olton and R.P. Kesner) 339-362 (Erlbaum, Hillsdale, NJ, 1990).

[16] Little, W.A. The existence of persistent states in the brain. *Math. Biosciences* **19**, 101-120 (1974).

[17] Hopfield, J.J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA* **79**, 2554-2558 (1982).

[18] Hopfield, J.J. Neurons with graded response properties have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA* **81**, 3088-3092 (1984).

[19] Kohonen,T. (1984) *Self-Organization and Associative Memory* (Springer-Verlag, New York, 1984).

[20] Marcus, C.M., and Westervelt, R.M. Dynamics of iterated-map neural networks. *Phys. Rev. A* **40**, 501-504 (1989).

[21] Brown, T.H., and Kairiss, E.W. *Neural Networks* (MIT Press, Cambridge, MA 1990).

[22] Gerstner, W. and van Hemmen, J.L. Associative memory in a network of 'spiking' neurons. *Network* **3**, 139-164 (1992).

[23] Gerstner, W. and van Hemmen, J.L. Coding and information processing in neural networks. In: *Models of neural networks III* (edited by E. Domany, J.L. van Hemmen, and K. Schulten) 1-93 (Springer, New York, 1995).

[24] Herrmann, M., Hertz, J.A., and Prügel-Bennett, A. Analysis of synfire chains. *Network* **6**, 403-414 (1995).

[25] Hopfield, J.J. and Herz, A.V.M. Rapid local synchronization of action potentials: toward computation with coupled integrate-and-fire neurons. *Proc. Natl. Acad. Sci. USA* **92**, 6655-6662 (1995).

[26] Maass, W. An efficient implementation of sigmoidal neural nets in temporal coding with noisy spiking. *NeuroCOLT Technical Report Series* NC-TR-96-031 (1996)

[27] Buhmann, J., and Schulten, K. Associative recognition and storage in a model network of physiological neurons. *Biol. Cybern.* **54**, 319-335 (1986).

[28] Shadlen, M.N., and Newsome, W.T. Noise, neural codes and cortical organization. *Current Opinion in Neurobiology* **4**, 569-579 (1994).

[29] Softky, W.R. Simple codes versus efficient codes. *Current Opinion in Neurobiology* **5**, 239-247 (1995).

[30] Koch, C., Rapp, M., and Segev, I. A brief history of time (constants). *Cerebral Cortex* **6**, 93-101 (1996).

[31] Markram, H., and Tsodyks, M. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* **382**, 807-810 (1996).