



ELSEVIER

Physica D 107 (1997) 392–400

PHYSICA D

Time for retrieval in recurrent associative memories

Alessandro Treves¹, Edmund T. Rolls*, Martin Simmen²

Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, UK

Abstract

Experimental evidence shows that certain types of visual information processing, such as face recognition, may be extremely rapid: in a few tens of milliseconds a neuron may yield most of the information about a visual stimulus that can ever be extracted from the response of that neuron. These data might be taken to indicate that there is insufficient time for recurrent or feedback processing. However, a novel analytical method allows the analysis of the dynamics of an associative memory network of integrate-and-fire neurons, laterally connected through realistically modelled synapses. The analysis, supported by preliminary simulations, indicates that the network may retrieve information from memory over a time determined mainly by excitatory synaptic conductance time constants, that is in times in the order of tens of milliseconds. The results thus show that the dynamics of recurrent processing is sufficiently rapid for it to contribute to processing within the short times observed in neurophysiological experiments.

Keywords: Feedback; Integrate-and-fire; Collective transitions

1. Introduction

Can one tell which connections are involved in processing information of a given sort in the cortex, by looking at the temporal course of the responses of individual neurons? These connections include those between neurons within a brain area which could implement attractor networks with recurrent connections; and feedforward connections from neurons in one area to neurons in the next area through a hierarchical processing scheme of connected areas. Thorpe and Imbert [1] have argued that the speed with which neurons in

the early visual system appear to be able to produce selective responses, e.g. when tested with different orientations of a stimulus, indicates that the processing is of a *feedforward* nature. They have recorded from orientation selective neurons that selectively respond to certain orientations of visual stimuli in the primary visual cortex of the monkey. The orientation tuning curves appear very similar if computed from the firing rate over a 300 ms period, or only over a 50 ms period after response onset. These results have been interpreted as ruling out the involvement of *recurrent feedback loops*, i.e. that a given neuron may alter its response over time, due to the firing of other cells which it itself influences (whether directly or polysynaptically). Examples of some classes of networks in which feedback effects may be important include autoassociative memories with recurrent excitatory connections, such as are found between cortical pyramidal cells, and also competitive networks with recurrent

* Corresponding author. Tel.: +44-1865-271348; fax: +44-1865-310447; e-mail: edmund.rolls@psy.ox.ac.uk.

¹ Present address: SISSA – Neuroscience Program, via Beirut 2-4, 34013 Trieste, Italy.

² Present address: Centre for Cognitive Science, University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, UK.

lateral inhibition [2]. Very similar interpretations to those of Thorpe and Imbert have been drawn by Oram and Perrett [3], from an analysis of the responses of ‘face’ cells in the temporal lobe visual cortex.

A series of experiments carried out in this lab [4,5] has reinforced the notion of rapid face processing in higher visual areas. The responses to face stimuli of neurons in the temporal cortex of macaques have been analysed using information-theoretical quantitative methods. The main finding is that a large proportion of the information about which face was shown is present if the firing rate is measured in a period as short as a few tens of ms taken near the onset of the neuronal response. For example, computing the rate over a 20 ms interval beginning 100 ms after stimulus onset (i.e. starting after the response latency of virtually all recorded cells) yielded on average 59% of the information obtained by computing the rate over a 500 ms interval beginning at the same time, while using a 50 ms interval yielded on average 78%. Further, the firing rate itself (a unidimensional measure of the response), when calculated over a long interval, contained nearly all the information (73%) that could be extracted from the same neurons by analysing the full temporal course of their responses (a highly multidimensional and potentially much richer measure). In addition, experiments on the neurophysiology of masking [6,7] have shown that neurons in the inferior temporal visual cortex fire for as little as 30 ms when a stimulus can just be recognized. Taken together, these findings suggest that in higher visual areas, in which cells are selective to complex patterns of visual stimulation, the processing involved, e.g. in object recognition is very rapid. The initial output, produced by each neuron shortly after receiving afferent inputs, is already information-rich, and is quickly refined within for example 50 ms of the initial response [4].

Do such experimental results really rule out, or severely limit, the role of feedback in networks in the brain during rapid object recognition? A feedback loop, it is to be noted, can be precisely defined only when referring to a specific cell A . It presupposes the emission of a sequence of spikes (σ) by a chain of synaptically connected neurons: $\sigma(A) \rightarrow \sigma(B) \rightarrow \dots \rightarrow \sigma(C) \rightarrow \sigma(A)$. The extent to which

feedback loops involving different chains determine the response of a whole population remains, however, an intrinsically ill-defined quantity. It is possible to resolve the ambiguity by using the more comprehensive notion of recurrent processing: the response of a local population of neurons is dependent on recurrent processing if it would be altered by removing lateral connections (those for which both the pre- and the post-synaptic cell belong to the population). The dependence can be quantified by quantifying the difference between the responses produced when such connections are present and when they are removed. What we ask, then, is whether recurrent processing may contribute, already within, e.g. 20 ms, to generate the experimentally observed selective responses.

Previous consideration of this issue relied, explicitly or implicitly, on considerably simplified dynamical models, in which either time was discretized into artificial time steps (‘cycles’) [8], or else the spiking, discontinuous nature of the dynamics was neglected altogether in favour of a description in terms of continuously varying firing rates [9]. Models of that kind reinforced the intuitive conclusion that cycling information through lateral connections necessarily involves using up precious time ‘cycles’; and that, hence, the observed speed indicates a streamlined processing mode, in which the response of each cell is determined by afferent inputs weighted by the strengths of synapses on afferent connections, and is quickly fed forward to the next cortical station.

We have developed a new analytical method for treating, in continuous time, the dynamics of a network of spike-emitting cells. The method considers a circumscribed population of neurons in which there are three sets of synaptic connections: afferents, efferents, and *recurrent collateral* ones. Here we present the results of detailed analyses [10] and preliminary simulations, and we apply them to address the question of recurrent processing in the visual system. This leads to new conclusions about the time constraints that the experimental data place on the type of processing.

The paper is organized as follows. Section 2 describes a formal model for understanding the temporal properties of recurrent neural networks. It also describes how a simulation of the model was

implemented. Section 3 describes the dynamic behaviour of the network both with the mean field analysis and with the simulation. Section 4 considers the implications of these results for understanding processing by networks with recurrent connectivity in the brain.

2. Methods

The analytical method we use is articulated as follows. First, we choose a suitable formal representation for individual units (neurons and synapses), that, though simplified, includes the salient features of their dynamical behaviour. We note that this model leads to behaviour similar to that of real neurons recorded in brain slices. Second, the network is defined in terms of its connectivity and synaptic strengths. Third, an analytical treatment of the behaviour of the network is described. Although a full dynamical treatment is complicated, we are able to analyse the approach to a type of asymptotic (long-time) behaviour. We then describe this asymptotic behaviour itself, in which the firing rate of each cell is stationary (provided that external inputs are kept, after an initial transient, constant), and also the transient modes, each with its own time constant, through which this behaviour is approached. A simple analytical expression yields the full spectrum of time constants for a given network, and, by allowing calculation of these time constants, explicitly relates them to the biophysical parameters introduced in defining the model.

Integrate-and-fire neurons. A simple point-like model neuron is considered, whose membrane potential, in between spikes, follows the usual *RC* equation [11]

$$C_i \frac{dV_i(t)}{dt} = g_i^0 (V_i^0 - V_i(t)) + g_i^K (V_i^K - V_i(t)) + \sum_{\alpha} g_i^{\alpha} (V_i^{\alpha} - V_i(t)). \quad (1)$$

The process of spike emission is faster than the time scales of interest, and hence it is not modelled in detail; rather, as the potential reaches a threshold value, a spike is added to the spike count, and the potential

itself is reset to a hyperpolarization potential, from which the evolution resumes as in Eq. (1).³ In Eq. (1) we have introduced an intrinsic potassium conductance in order to produce a simplified version of the phenomenon of rate adaptation. Its dynamics is simply described by assuming it to decay exponentially in between spikes and to be incremented by a fixed amount during each spike emission,

$$\frac{dg_i^K(t)}{dt} = -\frac{g_i^K(t)}{\tau^K} + \Delta g_i^K \sum_k \delta(t - t_{k,i}). \quad (2)$$

Synapses. Synaptic conductances, both intrinsic and on afferents, are modelled by a conventional α -function behaviour [13]. A spike emitted at time $t_{k,j\alpha}$ by a presynaptic cell $j\alpha$ activates, after a short interval (Δt) summarizing axonal and synaptic delays, a conductance on the postsynaptic membrane which then follows the equation

$$\tau^{\alpha} \frac{d^2 g^{\alpha}(t)}{dt^2} + 2 \frac{dg^{\alpha}(t)}{dt} = -\frac{g^{\alpha}(t)}{\tau^{\alpha}} + \Delta g^{\alpha} \sum_k \delta(t - \Delta t - t_{k,j\alpha}), \quad (3)$$

where τ^{α} is the corresponding synaptic time constant, characterizing the exponential return of the conductance to its closed state. In the simulations, synaptic conductances with instantaneous activation and exponential decay were used instead of the α -functions, and the delays Δt for pyramidal–pyramidal synapses were set randomly in the range 2–8 ms, and they were set to 1 ms for all other synapses.

To demonstrate that the single cell model is appropriate, Fig. 1 shows the response of a simulated pyramidal cell to a sequence of injected currents, illustrating the effect of firing-rate accommodation.

Network architectures. In principle, any architecture can be treated within this formalism, as long as cells of the same class (i.e. with similar properties)

³ While one may consider also an absolute refractory period (lasting e.g. 1–2 ms), we neglect that, for simplicity. In theory such a simplification makes possible infinite spiking rates, but in practice actual rates will be regulated, and at much lower levels, by other mechanisms such as inhibition [12], and the absolute refractory period will not therefore be influential.




Fig. 1. Response of simulated pyramidal cell (upper panel) to injected current (lower panel). Depolarization spikes are drawn in for clarity. Pyramidal parameters (used throughout paper): $C = 0.375$ nF, $g^0 = 25$ ns, $V^0 = -73$ mV, $V^{\text{thr}} = -53$ mV, $V^{\text{reset}} = -63$ mV, $V^K = -85$ mV, $V_E^E = 0$ mV, $V_E^I = -75$ mV, $\Delta g^K = 9.375$ ns (interneuron set similar, except $g^0 = 75$ ns and $\Delta g^K = 0$ nS).

are placed in a homogeneous group, and the network consists of a number N_C of such groups of cells. Uniform parameters should characterize each class, and the synapses between cells from any two classes. If the connectivity is chosen such that the existence of a synapse and the magnitude and time constant of its associated conductance depend only on the type of the pre- and post-synaptic cells, then the partition into classes may correspond to a physiological classification into cell types [14,15]. Alternatively, when considering, as we do here, an associative memory model in which information is stored in the detailed synaptic structure, one has to partition a physiologically homogeneous population of cells of one type into as many classes as required so that synaptic strengths (e.g. the magnitude of the relative conductances) are fully determined by the classes of the pre- and post-synaptic cells. This may result in many classes consisting of only few cells, or even one.

We simulated an ensemble of 8192 pyramidal-like cells coupled to 500 interneuron-like cells. The pyramidal cells are linked with each other and the interneurons through a sparse network of monosynaptic connections generated randomly. Connection probabilities were chosen broadly in line with hippocampal CA3 data [16], e.g. each pyramidal cell is innervated

by 10% of the others, each interneuron is innervated by on average 200 pyramidal cells and is presynaptic to a similar number. The strengths of excitation and inhibition depend on the connectivity and Δg values and are set by the parameters r_E^I , r_I^E , and r_E^E , where the subscript denotes cell type, and the superscript the type of synapses providing input to the cell. For example, r_E^I is $\sum \Delta g$ over the inhibitory (I) synapses of a typical (excitatory, E) pyramidal, normalized by the pyramidal g_E^0 . The pyramidal–pyramidal synapses encode a set of memory patterns, in this case forty random binary patterns each with 10% activity. This is achieved by first using the covariance rule to generate synaptic efficacies encoding these patterns [2], and then globally scaling them into positive Δg values such that the average per-cell sum equals $r_E^E g_E^0$.

Mean-field analysis. A mean-field description is obtained by summing up the equations describing the dynamics of individual units to get (fewer) equations that describe collective behaviour [17]. Thus grouping Eqs. (1) results in N_C functional equations describing the evolution in time of the fraction of cells of a particular class that at a given instant have a given membrane potential, while grouping Eqs. (4) results in $N_C \times N_C$ equations describing the dynamics of the summed conductance opened on the membrane of a cell of a particular class by all the cells of another given class. The system of mean-field equations has stationary solutions characterized by a constant firing rate for each cell class. As the neuronal current-to-frequency transfer function under stationary conditions is rather similar to a threshold-linear function, while each synaptic conductance is constant in time, the stationary solutions are essentially the same as can be obtained using much simpler, non-dynamical methods with threshold-linear units [18–21]. The advantage of the dynamical treatment, however, is that one can describe the approach to the asymptotic solutions, characterized by transient modes that decay with an exponential time dependence $e^{\lambda t}$. By linearizing the dynamical equations near a stable stationary solution and manipulating them algebraically, as explained in detail in [10], one obtains Eq. (4) yielding all the possible values of λ (they must all have a negative real part for the solution to be stable).

3. Results

3.1. Mean-field analysis

For any possible stationary asymptotic firing behaviour of the network of formal neurons described under ‘Methods’, there is an infinite set of transient modes, each associated with a complex time constant λ . The full λ spectrum for a given stationary state is obtained as the set of solutions of the complex equation

$$\Delta = |\mathbf{Q}(\lambda) - e^{\lambda \Delta t} \mathbf{1}| = 0, \quad (4)$$

where Δ denotes the determinant of a matrix of rank N_C (the number of classes), composed of the matrix $\mathbf{Q}(\lambda)$ and the identity matrix $\mathbf{1}$. The form of $\mathbf{Q}(\lambda)$ [10] is determined by the description chosen to model single units, i.e. neurons and synapses. Once that description, and all the parameters that accompany it, are specified, one has to solve Eq. (4) to find the time constants that govern the collective behaviour of the network. This last step may involve a straightforward numerical procedure, if N_C is small, or further analytical manipulation followed by numerical evaluation, if N_C is large or, in particular, tends to infinity [10]. In either case, the set of λ values obtained reflects the parameters describing cells and synapses, and thus explicitly links the single unit dynamics (the level at which the model attempts to capture relevant biophysical features) with network dynamics (where contact with neurophysiological experiment can be made).

The dependence of the spectrum on the underlying parameters is in general very complicated, but may be characterized as follows. The spectrum, plotted on the complex plane $\Re(\lambda)$, $\Im(\lambda)$, where both axes are measured in Hz, presents a gross, a fine and a hyperfine structure (see Fig. 2). The gross structure consists of those λ values which satisfy Eq. (4), whose real part is in the kHz range and beyond. These values correspond thus to fast time scales, and are determined by fast time parameters (e.g. 1 ms) such as the delay Δt . They are also very sensitive to global conditions like the detailed balance of excitation and inhibition, and in fact an instability associated with an imbalance of this sort may show up as one or more of the λ 's of the

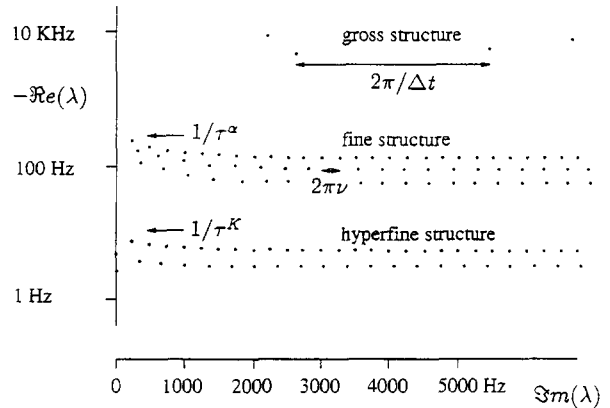


Fig. 2. Schematic representation of the spectrum of time constants characterizing collective dynamics, representing the dependence on the main model parameters, as found by numerically solving Eq. (4).

gross structure acquiring a positive real part. There are two reasons, however, for not focusing one's attention on the gross structure of the spectrum. One is that, inasmuch as fast phenomena characterising the real neural system have either been omitted altogether from the model (such as the Hodgkin–Huxley details of spike emission) or crudely simplified (such as unifying all axonal and synaptic delays into a unique delay parameter Δt), the model itself is not likely to be very informative about the fast dynamics of the real system. Second, in the presence of a stable stationary solution and of transients (those associated with the fine and hyperfine structures) lasting longer than, say 5 ms, the faster transients are not very significant anyway.

The fine and hyperfine structures both consist of orderly series of λ values satisfying Eq. (4), whose real parts are similar (within a series), while the imaginary parts take discrete values ranging from zero all the way to infinity, with approximate periodicity $\Delta \Im(\lambda) \simeq 2\pi\nu$, where ν is the firing rate of a particular group of cells. The (negative) real parts are complicated functions of all the parameters, but in broad terms they are determined by (and similar in value to) the inverse of the conductance time constants [10]. The fine and hyperfine structures differ in the magnitude of their real parts, i.e. in their time constants, in that relatively fast inactivating conductances (e.g., those associated with excitatory transmission via AMPA receptors) produce

the fine structure, while slower conductances (e.g. the intrinsic ones underlying firing-rate adaptation [22], or GABA_B inhibitory ones [23]) produce the hyperfine structure. The fine structure, which therefore corresponds to time scales intermediate between those of the gross and hyperfine structures, is perhaps the most interesting, as it covers the range in which the formal model best represents dynamical features of real neuronal networks.

In principle, having assigned a definite starting configuration for the model network, i.e. the distributions of membrane potentials of the cells and the degree of synaptic activation, at a given time, one should be able to follow in detail the dynamical evolution, thereby quantifying the relative importance of different time scales. In practice however, the explicit characterization of the modes with their spectrum of time constants, is only exact asymptotically close to a stationary state, and thus applies only to starting configurations infinitesimally close to the stationary state itself. For arbitrary initial conditions, a behaviour increasingly similar to that described by the modes and their time constants will progressively set in, provided the network approaches a stable stationary state of the type considered here.

What does the analysis of the formal model imply for the behaviour of real networks of neurons, e.g. in visual cortex? With the proviso mentioned in the last paragraph, the essential conclusions, for the model, are that (i) the approach to stable firing configurations is governed by a variety of time scales of widely different orders of magnitude; and (ii) an important subset of those time scales is determined by the inactivation times of synaptic conductances. These conclusions are expected to be valid, as well, for the dynamics of local cortical networks. With excitatory synapses between pyramidal cells [24], the time course of inactivation, even taking into account the spread due to the finite size of the postsynaptic cell (neglected by our treatment based on point-like neurons), may be as short as 5–10 ms. Therefore a substantial component in the ‘settling down’ of the distribution of firing activity into a stable distribution, under the influence of interactions mediated by lateral connections, may be very rapid, that is it may occur in little more than 5–10 ms.

3.2. Simulations

A cue consisting of a 20% fragment of one of the memory patterns (i.e., 0.1 bits/cell in information) was provided to the network, and its performance at retrieving the whole pattern was studied, primarily as a function of the excitatory feedback strength r_E^E , holding r_I^E and r_E^I at 20 and within the range 10–50, respectively. Two distinctive regimes were identified, termed here the *weak* feedback and *strong* feedback regimes.

Weak feedback. Here, the r_E^E values are such that the model pyramidal–pyramidal Δg is $O(1 \text{ ns})$. Cue-firing can generate activity in the non-cued pattern cells, but the pyramidal–pyramidal connections are too weak to allow persistence in attractor states after cue-removal. When the cue is active, synchronized oscillatory firing is seen and pattern retrieval is rapid but incomplete. As synchronized oscillatory activity entails periods of quiescence, we reasoned that this might hinder retrieval and therefore sought stationary states. To do this, the interneurons were injected with a bias current I_{ext} to prime them to respond faster to activity in the pyramidal net, to prevent the positive feedback in that net from leading to massive synchronous discharge. Fig. 3 shows the influence of priming, for $r_E^E = 40$. For neither level of priming does the retrieved information match the amount of information in the pattern ($\approx 0.47 \text{ bits/cell}$), although it reaches up to twice the level present in the cue (from 1 to 2 bits/cell).




Fig. 3. Lines in the upper panels denote population-averaged rates over 5 ms windows: dashed line – cued cells, dotted line – non-cued pattern cells, solid line – background cells, solid line + circles – interneurons.

In the oscillatory firing regime information retrieval is limited because the interneurons provide only feedback inhibition. Their late response to growth in the non-cued pattern cells allows the background activity to rise. This causes overlap between the rate distributions of background and pattern cells and therefore degrades retrieval. In the non-oscillatory regime, the early and continual interneuron activity helps not only by suppressing background activity but also reduces recruitment of non-cued pattern cells, slowing information retrieval. These results suggest that a better retrieval scenario would involve continual interneuron firing coupled with stronger inter-pyramidal feedback to sustain higher activity in the pattern cells.

Strong feedback. Here the r_E^E values are such that the model pyramidal–pyramidal Δg is 5–10 ns. Large model Δg 's are arguably more realistic, as the number of excitatory synapses per pyramidal cell in the simulation (≈ 819) is an order of magnitude down on the anatomical figure. Without interneuron priming, the cue activity spreads rapidly until within ≈ 15 ms every pyramidal is firing at an unrealistically high rate. With priming and a stimulus onset delay, there is an early phase of gradual growth in pattern cell activity, followed by ignition of firing throughout the entire network. Fig. 4 illustrates this for $r_E^E = 200$, $I_{\text{ext}} = 1$ nA. As shown, with the pattern cells firing at ≈ 50 Hz and the background cells quiescent, information retrieval occurs rapidly; but unfortunately this state is unstable. Extensive trials in nearby regions of parameter space find the '50 Hz' state either again building to network ignition or falling back to lower activity levels. The origins of such instability, and the conditions which avoid it, have recently been elucidated in related work by Battaglia and Treves [25]. The instability arises

because under the strong feedback necessary for retrieval to be completed by the excitatory recurrent connections, subtractive inhibition is insufficient to control the spread of excitation. It is necessary for inhibition to include a strong multiplicative component for the network to reach a genuine stable retrieval state.

Thus, large-scale simulations indicate, if tentatively, that it is possible, given small initial cues, to achieve partial retrieval of sparse patterns stored in a sparsely connected network of 'integrate-and-fire' neurons within ≈ 50 –100 ms of the cue presentation. This is a step towards the long-term goal of extracting predictions from such models about the dynamics of cortical memory retrieval. Full pattern retrieval was found to be hampered by the lack of a stable state with intermediate firing rates – this problem has been addressed by more recent work [25].

4. Discussion

What about the intuitive notion we referred to earlier, suggesting that feedback in such neuronal networks would be associated with much slower time scales? Such reasoning might proceed like this: (a) the stabilization of a firing configuration mediated by lateral connections is an inherently feedback effect; (b) for feedback to take place, the minimal requirement is that a cell contributes to the firing of another cell, then receives back a spike from that cell and fires again, using up in the order of two interspike intervals (one each at the current firing rate of each of the two cells); (c) the process is iterated a few times in order to lead to a steady activity distribution, taking up in the order of several hundreds of ms.

Such reasoning, which incidentally neglects both the spread in the latencies of activation of the cells in the network and, more importantly, their ongoing spontaneous activity, is a sequel of three incorrect steps. In fact (a) lateral connections may contribute much of their effect (mediated in our case by synapses modified through learning) well before genuine feedback occurs, in fact even before most cells have had the time to fire a single action potential in response




Fig. 4. Behaviour for $r_E^E = 200$, $I_{\text{ext}} = 1$ nA, cue onset at 10 ms, duration 50 ms.

to a new stimulus. For example, a spike emitted by cell A may immediately influence the firing of a cell B that happened to be close to threshold, while more time may be needed for the firing of B to have a feedback effect on the emission of a subsequent spike by A. (b) The influence, reciprocated or not, of a given cell onto another is only minor for the collective behaviour of the cell population, which is determined by the concerted action of the thousands of different inputs to each cell. (c) Reasoning in terms of time steps, or ‘iteration cycles’, is meaningless anyway, as the dynamics of the real system is in continuous time and, unless a special synchronization mechanism sets in, is asynchronous.

A formal model based on firmer assumptions leads, as we have seen, to different conclusions. These may be understood intuitively, if one considers that within a large population of cells, firing at low but non-zero rate, as occurs in cerebral cortex [26], incoming afferent inputs will find some cells that are already close to threshold. The slight modification in the firing times of those cells (brought forward or postponed depending on whether each cell is excited or inhibited by the afferents) will be rapidly transmitted through lateral connections, and (for example in associative networks) contribute to initiate a self-amplifying process of adjusting the firing pattern that would have been determined by external inputs alone. The process will proceed at a pace set by the time course of synaptic action, which in practice, in a situation in which transmission delays are negligible, means by the time constants of synaptic conductances. The fact that the speed of this ‘collateral effect’ [27] is essentially independent of prevailing firing rates is somewhat analogous to, and no more puzzling than, the propagation speed of elastic waves being distinct from the speed of motion of the individual particles carrying the wave.

Information-rich neuronal responses emerging at relatively early stages may therefore reflect not only feedforward but also, to a considerable degree, lateral interactions. As discussed above, lateral interactions may or may not involve, for any given cell, the propagation of spikes around feedback loops. The

sequential activation of most such loops will require more than 20 ms, but the population response will already within that time be dependent on the activation of recurrent connections. The ability of recurrent connections to contribute usefully to information processing is due to the fact that part of their effects occurs within a time determined primarily by the time constants of the synaptic conductances. Slower time scales, partly intrinsic and partly also associated with lateral interactions, such as those involving slow inhibition and those underlying adaptation processes, may produce, later on, further adjustments in the pattern of firing activity. Those adjustments may possibly, but not necessarily, increase the information content of neuronal responses. In addition, in the longer time scale of 50 or more ms, long range feedback loops, such as those which may be implemented by back-projections within and to the cerebral neocortex [28], may be important in a number of types of top-down influence on information processing.

It is interesting to note that our conclusions depend on the use of a formal network model in which we have introduced, in simplified form, important aspects of real neuronal and synaptic behaviour, such as conductance dynamics. This suggests that the evolution in time of neuronal activities in the brain may follow very different patterns, even qualitatively, from those that could be inferred from the behaviour of artificial dynamical systems in which those aspects are absent. The work described here does suggest that networks of neurons with dynamics of the type captured by integrate-and-fire neurons can interact sufficiently rapidly for autoassociative memory retrieval to occur in a biologically plausible and sufficiently short time. Similar time scales may apply to other networks such as competitive networks in which the competition is implemented through a population of inhibitory interneurons [2].

In conclusion, the formal analysis described here has shown that the dynamics of recurrent networks with conductance dynamics could operate very fast, in the order of the time constant of the synaptic conductances. The simulation showed that this can occur even when the recurrent network is retrieving information from a partial input. The formalism is not

restricted to that of recurrent attractor networks, but applies to networks with other architectures such as feedforward and competitive networks.

Acknowledgements

This research was supported by the UK Medical Research Council Interdisciplinary Research Centre in Brain and Behaviour (University of Oxford), by an EC BRAIN Initiative grant (Parallel Processing in Neuronal Networks), and by an EC Human Capital and Mobility grant (Computations in Neuronal Networks in the Brain). We also thank the EPCC (University of Edinburgh) for time on the CM-200 supercomputer.

References

- [1] S.J. Thorpe and M. Imbert, in: *Connectionism in Perspective*, eds. R. Pfeifer, Z. Schreter and F. Fogelman-Soulie (1989) pp. 63–92.
- [2] E.T. Rolls and A. Treves, *Neural Networks and Brain Function* (Oxford University Press, Oxford, 1997).
- [3] M.W. Oram and D.I. Perrett, *J. Neurophysiol.* 68 (1992) 70–84.
- [4] M.J. Tovee, E.T. Rolls, A. Treves and R.P. Bellis, *J. Neurophysiol.* 70 (1993) 640–654.
- [5] M.J. Tovee and E.T. Rolls, *Visual Cognition* 2 (1995) 35–58.
- [6] E.T. Rolls and M.J. Tovee, *Proc. Roy. Soc. B.* 257 (1994) 9–15.
- [7] E.T. Rolls, M.J. Tovee, D.G. Purcell, A.L. Stewart and P. Azzopardi, *Exp. Brain Res.* 101 (1994) 474–484.
- [8] D.E. Rumelhart and J.L. McClelland, *Parallel Distributed Processing* (MIT Press, Cambridge, MA, 1986).
- [9] H.R. Wilson and J.D. Cowan, *Biophys. J.* 12 (1974) 1.
- [10] A. Treves, *Network* 4 (1993) 259–284.
- [11] R.J. McGregor, *Neural and Brain Modelling* (Academic Press, San Diego, 1987).
- [12] L.F. Abbott, *Network* 2 (1991) 245–258.
- [13] T.H. Brown and D. Johnston, *J. Neurophysiol.* 50 (1983) 487–507.
- [14] D.A. McCormick, B.W. Connors, J.W. Lighthall and D.A. Prince, *J. Neurophysiol.* 54 (1985) 782–806.
- [15] A. Mason and A. Larkman, *J. Neurosci.* 10 (1990) 1415.
- [16] A. Treves and E.T. Rolls, *Hippocampus* 4 (1994) 374–391.
- [17] A.A. Frolov and A.V. Medvedev, *Biophysics* 31 (1986) 332–337.
- [18] H.K. Hartline and F. Ratliff, *J. Gen. Phys.* 40 (1957) 357.
- [19] K.P. Hadeler and D. Kuhn, *Biol. Cybern.* 56 (1987) 411.
- [20] A. Treves, *Phys. Rev. A* 42 (1990) 2418–2430.
- [21] A. Treves and E.T. Rolls, *Network* 2 (1991) 371–397.
- [22] D.A. Brown, B.H. Gähwiler, W.H. Griffith and J.V. Halliwell, *Progr. Brain Res.* 83 (1990) 141–160.
- [23] B.W. Connors, R.C. Malenka and L.R. Silva, *J. Physiol.* 406 (1988) 443.
- [24] A. Mason, A. Nicoll and K. Stratford, *J. Neurosci.* 11 (1991) 72–84.
- [25] F.P. Battaglia and A. Treves, *Soc. Neurosci. Abs.* 445.4 22 (1996) 1124; *Stable and Rapid recurrent processing in realistic autoassociative memories*, SISSA preprint 163/96/CNS (1996).
- [26] E.T. Rolls, *Phil. Trans. Roy. Soc. London B* 335 (1992) 11–21.
- [27] D. Marr, *Phil. Trans. Roy. Soc. London B* 262 (1971) 23–81.
- [28] E.T. Rolls, in: *Neural Models of Plasticity*, eds. J.H. Byrne and W.O. Berry (Academic Press, San Diego, 1989) pp. 240–265.
- [29] M. Mahowald and R. Douglas, *Nature* 354 (1991) 515–518.