

Neural circuits in silicon

Chris Diorio and Rajesh P. N. Rao

Studies of neurally inspired silicon circuits are showing how networks of neurons can select and multiply input signals. They may also provide alternative ways to build computers modelled on biology.

Animal brains form the centrepiece of nature's development of information-processing machines. In attributes such as adaptability and fault tolerance they are likely to remain unsurpassed by any human-built machine in the foreseeable future. Digital computers form the centrepiece of the past quarter-century's explosion in information technology. In terms of numerical computation, their capabilities far exceed those of animal brains. This dichotomy perplexes us. How do animal brains compute? And if silicon technology is so powerful, why can we not build thinking machines?

Carver Mead, a pioneer in integrated-circuit technology, quipped during a discussion of brains versus computers that "silicon doesn't know anything about bits". The implication is that nothing about silicon itself requires computers to be digital. On page 947 of this issue, Hahnloser *et al.*¹ take this reasoning to its logical conclusion: they have built a cortex-inspired silicon circuit that multiplies and selects features in its input, using a network of neuron-like elements. The theoretical basis for this work is not new: in 1996, Salinas and Abbott² reported computer simulations of a network of model neurons, demonstrating that the network could perform these exact tasks. What Hahnloser *et al.* have done is to reproduce the behaviour of such a network in silicon.

The networks studied by Salinas and Abbott and by Hahnloser *et al.* both use recurrent (or feedback) connections that are excitatory for connections between neighbouring neurons and inhibitory at larger distances (Fig. 1a). This pattern of local excitation and long-range inhibition is common in contemporary models of the cerebral cortex³⁻⁶. One attribute of this type of network is that, when a neuron at a given physical location receives an input, the network responds by activating both the stimulated neuron and a cluster of neurons around it (Fig. 1b). Moreover, when the background input to all neurons is increased systematically, this cluster of activity is multiplied by a gain factor that is a linear function of the background input (Fig. 1b). These responses are intriguing from a neurophysiological perspective. Say, for example, we equate the background input to a motor signal representing eye position, then these multiplicative responses are similar to the modulation of neuronal

responses by eye position observed in the visual cortex⁷.

A second important feature of this type of recurrent network is its ability to select

a single stimulus when presented with several competing stimuli. In this case the network exhibits nonlinear behaviour, selecting the strongest of the competing stimuli and

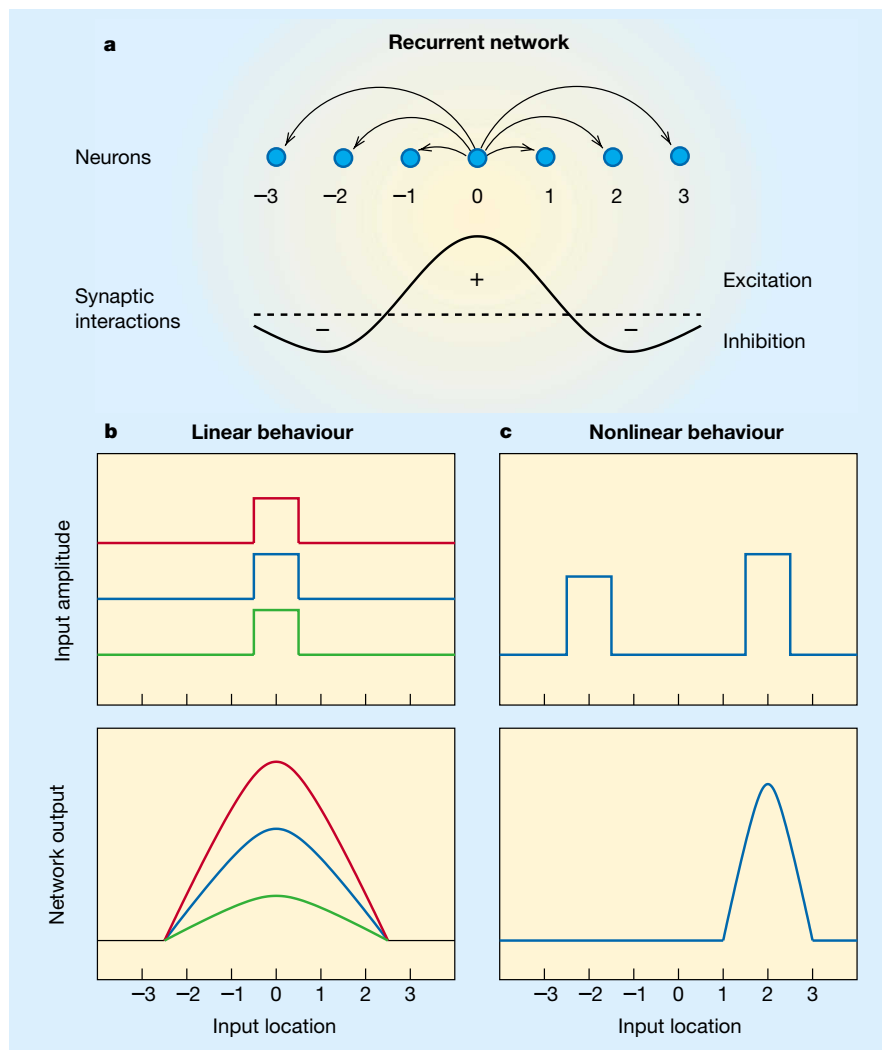


Figure 1 Linear and nonlinear behaviour in a silicon circuit inspired by neurobiology.

a, The neuronal network studied by Salinas and Abbott², and built in silicon by Hahnloser *et al.*¹.

Each neuron excites its nearest neighbours, and inhibits neurons further away from it. The artificial neurons, in both cases, use a continuously varying signal to model the firing rate of a cortical neuron. In their numerical simulations, Salinas and Abbott model the strength of synaptic connections from a neuron to its neighbours as a 'Mexican hat' function (shown below the network), whereas Hahnloser *et al.* approximate this function in their silicon circuit. b, Example of multiplicative responses (that is, linear gain modulation) in the network. Given three input pulses at the same location but at different background amplitudes (green, blue and red), the network multiplies the output by a gain factor that is a linear function of the background input. c, Example of stimulus selection in the network. Given input pulses at two different locations, the network selects the location with the stronger input, suppressing the location with the weaker input. This competitive, nonlinear behaviour is the result of recurrent excitatory and inhibitory interactions among neurons.

suppressing the weaker ones (Fig. 1c). Such behaviour can be regarded as a simplified form of sensory attention, whereby the network selects the stimulus location based on stimulus strength. From a neurophysiological perspective, the selection of a single target and the suppression of distractors is important, for example, in programming arm or eye movements.

The network's ability to switch between linear and nonlinear behaviour, based on its input, is very different from that of standard electronics. Engineers usually require separate analogue and digital circuits to carry out linear amplification and nonlinear selection, respectively. Hahnloser *et al.* explain their hybrid analogue–digital circuit in terms of the set of neurons that are active in steady state, and a gain that depends only on the identities of the active neurons and not on their analogue responses. Behaviour that derives from a common set of active neurons is linear in the input, whereas behaviour that derives from a comparison among different sets of active neurons is nonlinear in the input.

Hahnloser *et al.* extend previous studies of this class of recurrent networks by analysing the stability of the dynamical equations that model the network, and show that there are inviolate constraints on the allowed network states. In particular, they show that if no single input can activate a set of neurons (the set cannot form a memory), then no input can activate a supergroup of these same neurons (no supergroup can form a memory). A limitation of the study, however, is the requirement that synaptic connections be symmetric — that is, the strength of the connection from neuron A to neuron B must be the same as that from B to A. This assumption allows them to prove network stability, but is difficult to justify neurobiologically. A fruitful area for future research, therefore, is investigation of recurrent cortical networks with non-symmetric synaptic connections (see, for example, refs 8–10 and Supplementary Information¹).

What do Hahnloser *et al.*, and others like them, hope to accomplish by building silicon circuits modelled on biology? First, they can learn how to map neuronal 'primitives' (such as neurons and synapses) onto silicon, and then how to compute using these primitives (see, for example, refs 10–12). Second, they can investigate how physical and technological limits, such as wire density, signal delays and noise, constrain neuronal computation. And third, they can learn about alternative models of computation. Biology provides examples of non-digital computing machines that are incredibly space- and energy-efficient, and that excel at finding good solutions to ill-posed problems. Scientists may eventually decipher all of nature's electrochemical circuits, but the work of Hahnloser *et al.* demonstrates that we

already know enough to begin building integrated circuits that compute like biology. ■

Chris Diorio is in the Department of Computer Science and Engineering, University of Washington, Box 352350, 114 Sieg Hall, Seattle, Washington 98195-2350, USA.

e-mail: diorio@cs.washington.edu

Rajesh P. N. Rao is at the Sloan Center for Theoretical Neurobiology and the Computational Neurobiology Laboratory, Salk Institute for Biological Studies, 10010 N. Torrey Pines Road, California 92037, USA.

e-mail: rao@salk.edu

1. Hahnloser, R. H. R., Sarpeshkar, R., Mahowald, M. A., Douglas, R. J. & Seung, H. S. *Nature* **405**, 947–951 (2000).
2. Salinas, E. & Abbott, L. F. *Proc. Natl Acad. Sci. USA* **93**, 11956–11961 (1996).

3. Ben-Yishai, R., Bar-Or, R. L. & Sompolinsky H. *Proc. Natl Acad. Sci. USA* **92**, 3844–3848 (1995).
4. Somers, D. C., Nelson, S. B. & Sur, M. *J. Neurosci.* **8**, 5448–5465 (1995).
5. Douglas, R. J. *et al. Science* **269**, 981–985 (1995).
6. Stemmler, M., Usher, M. & Niebur, E. *Science* **269**, 1877–1880 (1995).
7. Andersen, R. A., Essick, G. K. & Siegel, R. M. *Science* **230**, 456–458 (1985).
8. Abbott, L. F. & Blum, K. I. *Cereb. Cortex* **6**, 406–416 (1996).
9. Rao, R. P. N. & Sejnowski, T. J. in *Advances in Neural Information Processing Systems* Vol. 12 (eds Solla, S. A., Leen, T. K. & Müller, K.-R.) 164–170 (MIT Press, Cambridge, MA, 2000).
10. Westerman, W. C., Northmore, D. P. M. & Elias, J. G. *Analog Integrated Circuits Signal Process.* **18**, 141–152 (1999).
11. Diorio, C., Hasler, P., Minch, B. A. & Mead, C. *IEEE Trans. Electron Devices* **44**, 2281–2289 (1997).
12. Mead, C. *Analog VLSI and Neural Systems* (Addison-Wesley, Reading, MA, 1989).

Neurobiology

Self-repair in the brain

Anders Björklund and Olle Lindvall

In tissues that can repair themselves, such as skin and liver, dead cells can be replaced either by the proliferation of nearby cells or by the activation of resident stem cells — undifferentiated cells with the potential to generate many different cell types. The brain apparently lacks this regenerative capacity, making it particularly vulnerable to injury or disease. Cells with stem-cell-like properties are likely to occur throughout the adult central nervous system, but they normally give rise to neurons in only a few, restricted areas. These cells might represent a dormant capacity for neuronal repair, if they could be mobilized to generate new functional neurons in response to injury. On page 951 of this issue, Magavi, Leavitt and Macklis¹ report the first intriguing evidence that this may be possible.

In the adult brain, the generation of new neurons (neurogenesis) occurs in just two regions². The first is the subventricular zone (SVZ) in the wall of the lateral ventricle. Here, new interneurons are generated for the olfactory bulb (Fig. 1a), which is involved in sensing odours. The second is the subgranular zone of the dentate gyrus, which gives rise to a different type of neuron, the granule cell. In these areas, there is seemingly a continuous turnover of interneurons and granule cells, implying that the newborn neurons replace dying cells. But does neurogenesis actually help to replenish damaged neuronal circuits? There is some evidence that damage to granule cells can trigger the increased proliferation and recruitment of new granule cells from resident progenitors³. And insults — such as seizures and inadequate blood supply to the cerebrum — that cause cells in the hippocampus to die are accompanied by increased neurogenesis in the dentate subgranular zone^{4–6}. But as yet there is no direct

evidence that the degenerated neurons are actually replaced by the new ones.

In other parts of the central nervous system, neural progenitors contribute to the ongoing formation of new non-neuronal cells called astrocytes and oligodendrocytes, and may participate in the reaction of these so-called glial cells to injury, as well as in scar formation^{7,8}. Other progenitors appear to remain undifferentiated, or die^{7,8}. A low level of neurogenesis has been described in parts of the intact neocortex of adult monkeys⁹, but the newly generated neuron-like cells appear to survive only transiently, and may not differentiate into fully functional neurons¹⁰.

In line with previous findings^{11,12}, Magavi *et al.*¹ have observed the continuous formation of new cells in the intact neocortex of the adult mouse. These proliferating cells were distinguished by labelling them with bromodeoxyuridine, a thymidine-base analogue that is incorporated into the new DNA formed during DNA replication in dividing cells. None of these newly formed cells, however, expressed any marker proteins characteristic of neurons.

To investigate the effects of neuronal death in the neocortex, Magavi *et al.* destroyed a subset of pyramidal neurons that project from the neocortex to the thalamus, which is a major subcortical relay station involved in the control of cortical function. Their technique resulted in the slow death by apoptosis of the targeted cells only, without affecting the surrounding cortical tissue. The number of new cells formed in the two weeks after the lesion was similar in control and experimental mice. However, about 1–2% of the newly formed cells in the damaged neocortex (about 50–100 cells per cubic millimetre) expressed neuronal markers.