Pointing the way toward target selection

L. F. Abbott

Hahnloser and colleagues propose a computational model in which top-down inputs, such as attention, act by modulating the strength of recurrent connectivity in a neuronal circuit.

Reaching for an object or shifting our eyes in a new direction requires the selection of a single target from what is typically a complex visual scene. To select a target, we must combine our knowledge of the visual scene with information about our interests and intentions. A paper by Hahnloser, Douglas, Mahowald and Hepp in this issue[1] addresses how this might be done by a neural circuit. The authors use a computational model to explore the idea that attentional inputs exert their influence by modifying the way that neurons in a target-selection network are interconnected. The idea that attention and other forms of top-down control might modulate the effective connectivity of a neural circuit is a powerful one that is likely to have a number of applications in the study of cognitive processing.

The fibers and synapses that interconnect different regions along a sensory pathway can be divided into three broad classes: feedforward, recurrent and top-down[2]. Feedforward connections carry input to a given region from areas that lie earlier along the pathway. Recurrent synapses interconnect neurons within a given region, and top-down connections carry information back from higher cognitive areas. Recurrent connections in a network can strongly affect how neurons respond to the feedforward input they receive. For example, a number of modeling studies have explored the idea that recurrent connections have a dominant role in shaping the response characteristics of neurons in primary visual cortex[3–7]. Modifying the nature of the recurrent connectivity in a neural circuit can have a dramatic effect on response selectivity and thus provides a powerful mechanism for controlling and modulating network function.

Hahnloser and colleagues[1] propose a way in which top-down inputs can modulate the recurrent connectivity of a neural circuit. The basic idea is that the recurrent feedback in the network is channeled through a relatively small number of neurons called pointer neurons (**Fig. 1**). As long as the pointer neurons are active, they provide a feedback pathway between the neurons of the network. This produces effective recurrent couplings with strengths given by the matrix product of the strengths of the connections from the network to the pointer neurons and from the pointer neurons back to the network. If, however, the pointer neurons are silenced by inhibition or otherwise modulated, the recurrent connectivity of the network will change.

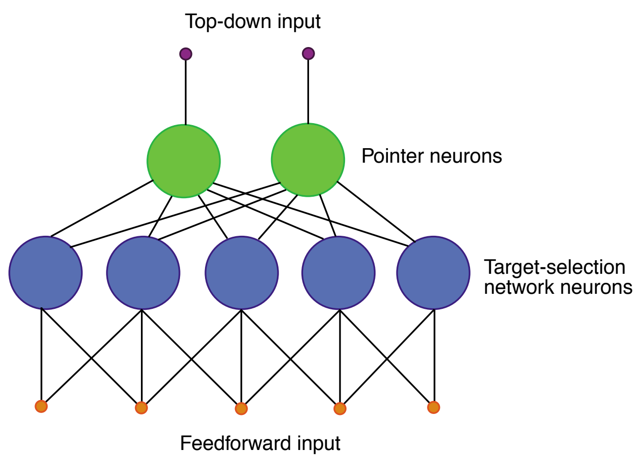

Fig. 1. The pointer circuit architecture. Feedforward connections carry inputs to the network. Recurrent connections between network neurons are funneled through pointer neurons that receive top-down input reflecting the locus of attention. Global inhibition that is present in the model of Hahnloser and colleagues[1] is not indicated in this figure.

The authors apply the idea of pointer neurons to a network that selects a target from the feedforward visual input it receives. The neurons in the target-selection network are arranged in a spatial map. A selected target is represented within this map by a 'hill' of activity in a localized population of neurons (**Fig. 2a** and **b**). The cluster of active neurons can be centered at different places on the map to represent different target locations. The selection of a unique desired target requires this network to generate a single hill of activity even if the input is complex (as in **Fig. 2c**), and this hill must correspond to a region of the visual field that is of interest.

As was known previously, the selection of a single target from a complex input can be accomplished through the appropriate choice of recurrent connectivity within the target-selection network. This makes use of a characteristic property of recurrent networks—their ability to support only particular stereotyped patterns of activity[3,8,9]. If the recurrent interactions within a network map are of the appropriate form, the network will only support a single hill of activity even if its feedforward input is complex[8,9]. This captures the basic 'winner-take-all' nature of the target selection process and is in agreement with experimental findings, for example, concerning the sparse representation of salient visual features in the parietal cortical area LIP[10]. The selection of a single target requires that inhibition of the network is strong enough to suppress responses to extraneous input, and that excitation is localized so that it can overcome inhibition within the hill of activity representing the selected target. In a recurrent network without top-down control, the single hill of activity in the target-selection network will typically be centered at the location of the strongest input (**Fig. 2a**). This would correspond to inevitably choosing the brightest image in a visual scene as a target.

The pointer architecture proposed by Hahnloser and colleagues[1] allows top-down influences of attention to overcome the tendency of the target-selection network to choose the strongest input. Attention enters the target-selection process and affects the responses of the network

L.F. Abbott is at the Volen Center and Department of Biology, Brandeis University, Waltham, Massachusetts 02454, USA. e–mail: abbott@brandeis.edu

Fig. 2. Target selection by a network map. The lower panel (**c**) indicates the level of input as a function of location for a simulated visual scene. Each bump represents a different object. The upper panels (**a, b**) show the level of neuronal activity in a target-selection network as a function of the position of the neurons within the map. Hills of activity represent selected target locations. In (**a**) the hill of activity corresponds to the location of the strongest input. Attentional affects implemented in the model of Hahnloser and colleagues allow the hill of activity to select and represent one of the weaker inputs, as in (**b**). This figure is a schematic and not the result of a network simulation.

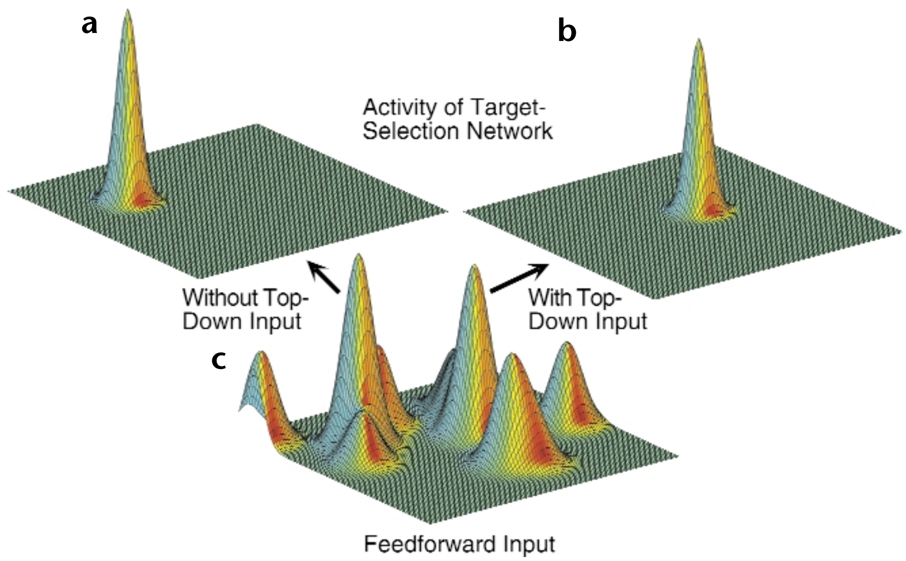

neurons through top-down input to the pointer neurons (**Fig. 1**). This can bias the network so that weaker inputs located in regions where attention is focused can determine the location represented by the hill of activity, even if stronger inputs are present elsewhere (**Fig. 2b**). This corresponds to choosing a general region of interest and then allowing the visual system to select a target within this region.

Recurrent networks can perform a number of other computations of relevance to sensory processing. For example, if the recurrent connections are strong enough, a particular hill of activity can be maintained even after the structured visual input is removed[3,8]. This provides a potential mechanism for short-term memory of the selected target location[11]. Hahnloser and colleagues[1] show that top-down inputs to the pointer neurons can bias or shift this remembered location, so that attention can affect sustained activity even when the feedforward input that originally produced it is no longer present.

Obviously, the critical question raised by this work is whether anything analogous to a pointer neuron actually exists. If the model is taken literally, these neurons may be very rare and thus would be difficult to find. However, the small number of pointer neurons in the model may not be a general feature of the pointer architecture. Pointer neurons might be distinguished in recordings because they are significantly less selective than the network neurons they interconnect, a result of receiving convergent input from neurons with a broad range of selectivities. The massive convergence and divergence of recurrent connections at a pointer neuron should also provide a distinctive anatomical signature.

If real neurons analogous to the pointer neurons do exist, they would be potent sites for all sorts of neuromodulation. Modulation of pointer neurons would be an effective way of changing the functional properties of a network because the pointer architecture allows the modulation of a small number of neurons to mimic the effect of modifying a large number of recurrent synapses. Hahnloser and colleagues only studied excitation or inhibition of the pointer neurons, and in the inhibitory case alone did the nonlinear features of the model come into play. Excitation of a pointer neuron simply enhances activity within a certain region of the network map. When a pointer neuron is inhibited from responding, an entire block of recurrent connections is removed from the circuit, which can have a more dramatic effect on neuronal selectivity. It would be interesting to explore a wider variety of pointer-neuron modulations. For example attention can modify the gain of a neuron[12–14], and gain modulation of pointer neurons would provide a graded way of changing network connectivity.

If attention and intention are to influence our perceptions, top-down connections must exert considerable control over sensory processing pathways. The hypothesis that they do this through their effects on recurrent interactions is a powerful one that deserves further exploration using both theoretical and experimental approaches.

1. Hahnloser, R., Douglas, R. J., Mahowald, M. & Hepp, K. *Nat. Neurosci.* **2**, 746–752 (1999).
2. White, E. L. *Cortical Circuits: Synaptic Organization of the Cerebral Cortex* (Birkauser, Boston, 1989).
3. Ben-Yishai, R., Bar-Or, L. & Sompolinsky, H. *Proc. Natl. Acad. Sci. USA* **92**, 3844–3848 (1995).
4. Somers, D. C., Nelson, S. B. & Sur, M. *J. Neurosci.* **15**, 5448–5465 (1995).
5. Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A. & Suarez, H. H. *Science* **269**, 981–985 (1995).
6. Maex, R. & Orban, G. A. *J. Neurophysiol.* **75**, 1515–1545 (1996).
7. Chance, F. S., Nelson, S. B. & Abbott, L. F. *Nat. Neurosci.* **2**, 277–282 (1999).
8. Salinas, E. & Abbott, L. F. *Proc. Natl. Acad. Sci. USA* **93**, 11956–11961 (1996).
9. Pouget, A., Zhang, K., Deneve, S. & Latham, P. *Neural Comput.* **10**, 373–401 (1990).
10. Gottlieb, J. P., Kusunoki, M. & Goldberg, M. E. *Nature* **391**, 481–484 (1998).
11. Camperi, M. & Wang, X.-J. *J. Comput. Neurosci.* **5**, 383–405 (1998).
12. Connor, C. E., Gallant, J. L., Preddie, D. C. & Van Essen, D. C. *J. Neurophysiol.* **75**, 1306–1308 (1996).
13. Connor, C. E., Preddie, D. C., Gallant, J. L. & Van Essen, D. C. *J. Neurosci.* **17**, 3201–3214 (1997).
14. McAdams, C. J. & Maunsell, J. H. *J. Neurosci.* **19**, 431–441 (1999).