Cell Assemblies for Rapid Recognition and Disambiguation

Christian R. Huyck
Middlesex University
c.huyck@mdx.ac.uk

Abstract

This document discusses reverberating circuits of neurons or Cell Assemblies (CAs) derived from Hebb's [2] initial proposal. It shows how CAs can quickly categorise an input and make a quick decision when presented with ambiguous data.

CAs are a widely accepted theoretical construct in cognitive neuroscience. This document also references some initial computational modelling of CAs.

CAs are thus an ideal mechanism for further research into both computational and cognitive neural models. If humans use CAs as a basis of thought, then studying how biological systems use CAs will provide information for computational models. The reverse is also true; computational modelling can direct our research activity in biological systems.

1 Introduction and Background

There are many possible connectionist architectures. The most successful architecture is implemented in slow moving biological circuitry. It exists in trillions of organisms, it can learn, be used for language and can even exhibit consciousness. Animal brains are the apex of intelligence. We have a pretty good understanding of how neurons work (by no means complete), but we have much less understanding of how neurons work together.

An average human neuron fires for a brief time, and then fatigues. If such a neuron were to represent a concept, we could only think of that concept for a brief time. Since we can think of a given concept for at least several seconds, a single neuron cannot represent an entire concept. A concept must be represented by a collection of neurons. How might such a concept be represented?

D. O. Hebb proposed a solution to this problem [2], a reverberating circuit of cells. Hebb called this reverberating circuit a *Cell Assembly* (CA). Some physiological evidence exists (e.g. [1]).

While there has been a fair amount of work in psychology [6] and neuro-biology to show the existence of CAs, there has been little computational modelling of CAs [3, 4, 5]. The computational modeling has either been from a more abstract level than neurons [5], or has been of a very limited nature [3, 4].

2 Quick Activation of CAs

The theory of CAs states that a CA is quickly activated [5]. This shows that slow neurons can quickly recognise an object.

Experiments using simple computational models of CAs shows this to be the case [4]. This model shows that activation of 10 neurons can lead to activation of the entire CA of 200 neurons in the next time step. A time step is meant to model 10 ms. of real neural activity. This particular experiment is probably not viable from a neurological standpoint as the size of the CA is very small. However, it does show how quickly a few neurons can lead to activation of many more neurons.

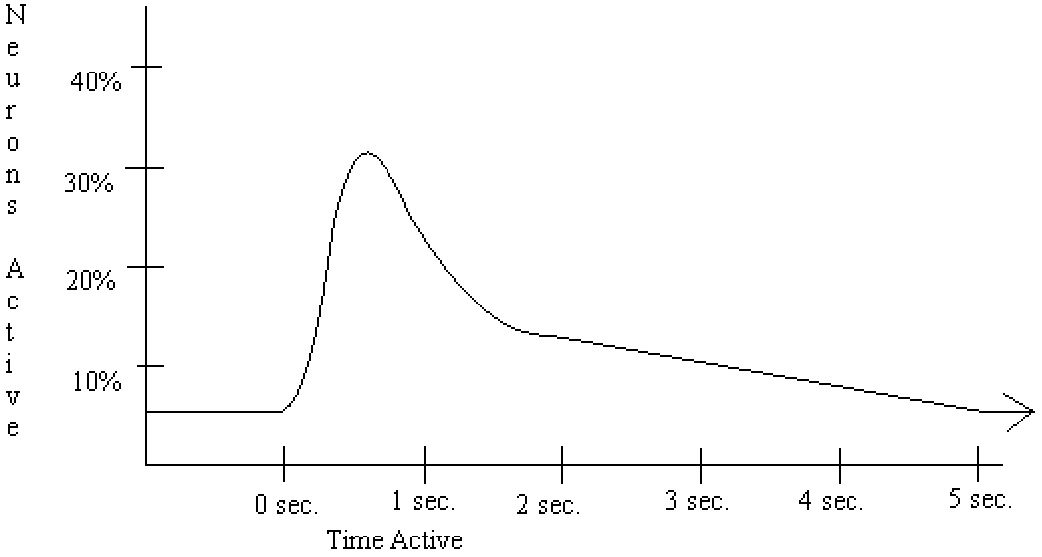

Figure 1.

A real CA would probably cross several brain areas. However, it would only take a few time steps (tens to hundreds of miliseconds) to activate the entire assembly. Highly parallel activation compensates for slow neural speed. Practically, it will be easy to get large numbers of neurons activated within the speeds that are needed for normal real-time performance.

Figure 1 shows the activation curve of a CA. At time 0, a stimulus is presented to the network. Rapidly, a large percentage of the neurons in the CA become activated and this is associated with recognition. Stimulus may be removed or remain present. Some neurons in the CA are active for quite some time meaning the CA can spread information, strengthen itself (via the Hebbian learning rule) and remain in short-term memory. Note that no neuron remains active for the whole period; they become active, fatigue, recover and are reactivated by other neurons in the CA.

3 Quick Decision on Ambiguous Data

A human brain must have millions of CAs, but a model network is much smaller. A given net may have several CAs in it, but only one or a few should be active at a given time. What happens when a net is presented with ambiguous data?

When ambiguous data is presented to a net, some of the neurons from two separate CAs will be activated. These will tend to activate other neurons in their own CAs. Since the two CAs have never been active together and reside in the same area of a distance-biased network, they will tend to inhibit each other. Thus a competition ensues between the two ambiguous CAs. One CA wins, and the data is recognised as an instance of that type of object.

When the net is presented with ambigous data it does take slightly longer to recognise the object. This is due to the extra time needed for competition. However, it should only take a few more cycles to recognise the object, and thus is almost as fast as recognising an object from unambiguous data.

Figure 2 shows the activation pattern of two CAs. Both get an initial bust of activating neurons. Some of these neurons are inhibitory neurons and inhibit neurons from other CAS this inhibiting the other CA. The CA represented by the bold line loses the competition and returns to its base firing rate.

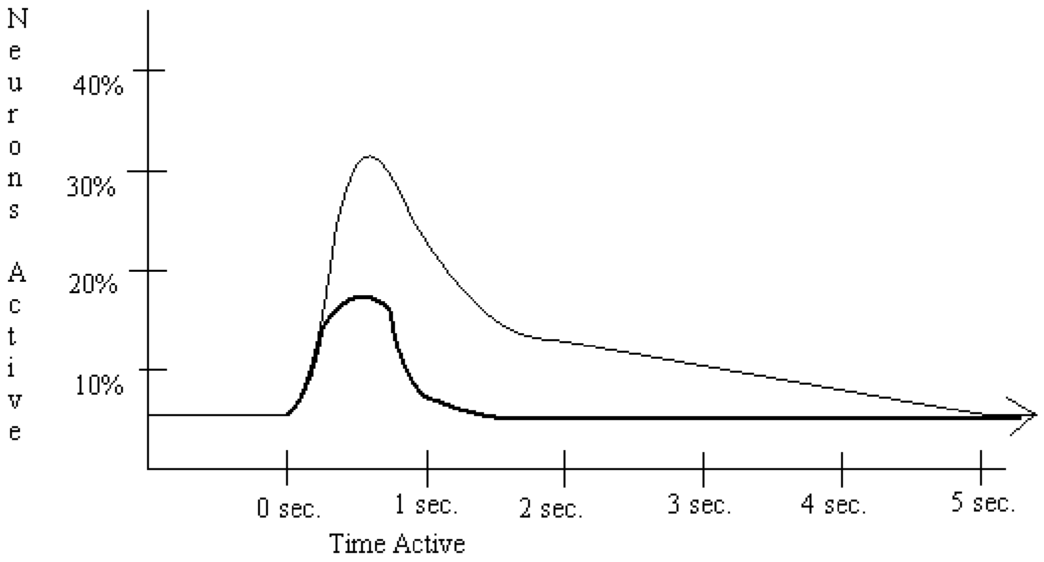
Neurons Active
40%
30%
20%
10%
0 sec.
1 sec.
2 sec.
3 sec.
4 sec.
5 sec.
Time Active

will be simplified as modelling a single neuron is very complex. Perhaps we can simplify the model by simply modelling changes in neural activation in discrete steps and learning by changes of stength. This is not as robust as modelling neural activation continuously with ionic transfer between neurons, and modelling changes of strength by changes in axonal radius, myelination and geometric properties of the synaptic cleft. This simplification makes the model run faster which should enable us to model the large number of neurons needed for CAs. Still it is a simplification and, as such, may miss important data.

It is more difficult to say what CAs can tell us about building new computational neural architectures. CAs are models for concept storage. They can solve the symbol grounding problem. Perhaps the most interesting question is in what way can CAs work together.

At heart, CAs are a model for data storage. They have several advantages that could be transfered to other models. They remain active for a certain period of time. They are a model for both long and short-term memory. They can be used to choose between equally likely solutions. It is the reverberating activity that makes this model different from most existing ANN models.

All neurons are not part of CAs. Neurons in the eye, send signals back to the Laterate Geniculate Nucleus and never get activated by other neurons. All other neurons may be part of CAs. (The present data is unclear.) Those neurons coming from the eye (and other sensory interfaces) are relatively well understood. Using those 'understood' neurons as inputs to a CA model may enable us to understand early feature detection. Thus we can start programming systems that work on real-world features.

CAs as the basis for combining CAs (cognitive maps). Activation for a CA can come from the environment or from other CAs. They can be used as models of associative memory, memory of sequences, and hierarchical memory. They might be used for rule based systems and other, traditionally symbolic, activities.

CAs are not modular though network topologies may be varied. Despite our assigning brain areas, the brain is also not modular. Each area has billions of connections to and from it. CAs can act as the basis of a wide range of architectures for solving a wide variety of problems in a massively parallel way.

Using CAs as a basis for more complex phenomena puts us on reasonably solid ground for future exploration of how the brain works at more complex levels. This in turn leaves us with a better understanding of intelligence.

References

[1] Abeles, M., H. Bergman, E. Margalit, and E. Vaadia. (1993) Spatiotemporal Firing Patterns in the Fronatal Cortex of Behaving Monkeys. *Journal of Neurophysiology* 70(4):1629-38

[2] Hebb, D.O. (1949) The Organization of Behavior. John Wiley and Sons, New York.

[3] Hetherington, P. A., and M. Shapiro. (1993) Simulating Hebb cell assemblies: the necessity for partitioned dendritic trees and a post-net-pre LTD rule. *Network: Computation in Neural Systems* 4:135-153

[4] Huyck, Christian R. (accepted). Modelling Cell Assemblies. In *Proceedings of the International Conference of Artificial Intelligence* Las Vegas, Nevada

[5] Kaplan, S., M. Sontag, and E. Chown. (1991) Tracing recurrent activity in cognitive elements(TRACE): A model of temporal dynamics in a cell assembly. *Connection Science* 3:179-206

[6] Pulvermuller, F. (1999) Words in the brain's language. In *Behavioral and Brain Sciences* 22 pp. 253-336.