Models of Cell Assembly Decay.

P.J. Passmore and C. Huyck
School of Computing Science
Middlesex University
London
UK.
NW4 4BT.
Email: p.passmore@mdx.ac.uk; c.huyck@mdx.ac.uk

***Abstract*—Hebb considered the cell assembly to represent a concept in the brain and thus to be an underlying construct of human thought. He proposed that the cell assembly is a connected group of neurons whose pattern of firing is such that a reverberatory activity persists after the removal of a stimulus. Once a cell assembly is activated something must eventually cause it to decay. Clearly thoughts have to be extinguished to make way for others, the question is how. Various suggestions have been made concerning mechanisms that could cause an assembly to decay in the long term including inhibition by other assemblies and passive fatigue. In this paper two classes of models are used to implement this decay, the first is based on building cell assemblies with specific weights and connections that have a linear decay. The second class is based on manipulating variables within a cell assembly model, creating long term fatigue or activation decay. This class of models may be more biologically plausible than the first, and can produce the expected temporal dynamics in the presence of an ambiguous stimulus. However neither class can yet produce the correct prolongation of activation when the stimulus is re-presented.**

I. Introduction

A cell assembly is an interconnected set of neurons which when stimulated form a reverberating circuit of activity. The cell assembly is a construct that was proposed by Hebb [6] to serve as a concept or unit of thought; since the initial proposal this hypothesis has been supported by many studies (e.g. [1], [15]). A central tenet of Hebb's approach is that cells that are co-located will form excitatory connections under repetitive neural activity, cells that "wire together fire together". However if such a process is left unchecked activation would lead to the uncontrolled spread of excitation. This led Milner [13] to later add inhibition to the model to counter this effect. By this model, storing of information in memory is seen to be achieved by the formation of cell assemblies. Hebb considered that long term memory was formed by changes in synaptic strength resulting from his excitatory rule.

In a paper on the temporal dynamics of cell assemblies, Kaplan et. al. [10] suggested that the activation of a cell assembly comes in a series of phases, giving it the power to handle a wide range of psychological data. They add fatigue and short term connection strength to the model of a cell assembly. A control on the amount of activity is modelled by fatigue, the tendency for units to drop out over time. Furthermore the authors suggest that local excitation will produce local inhibition and make a cell assembly vulnerable to competition from other cell assemblies. This then gives a particular temporal activity profile to a cell assembly. As excitation builds up activity increases as the cell assembly ignites, activity peaks and a longer period of slowly decreasing activity ensues as fatigue, local inhibition and competition have effect, until the cell assembly is extinguished. This gives the distinctive curve shown in figure 1.

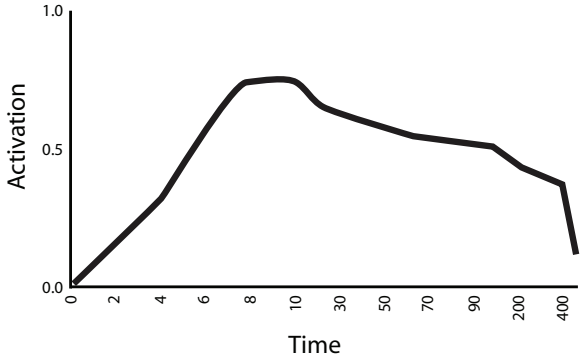

Fig. 1. Activity in a cell assembly over time. (After [10].)

While the above is one account for how excitation is extinguished it is not the only one. Burgess and Hitch (1999) [3] produced a connectionist model of the phonological loop (part of Baddley's 1986 [2] model of short term or working memory) in which connection weights are assumed to decay over time. An alternative model of working memory [14] proposed decay of a node's activation over time. Distributed models of working memory (eg: [4]) typically ascribe forgetting to interference, while [5] proposed self inhibition as a method of extinguishing "attentional shrouds". Beyond these possible causes for cell assembly activity decay are possible long term leak or fatigue processes.

In the ACT-R rule-based model, memory chunks' activation decay as a function of usage history and time according to a function which is intended to track the log odds that an item will need to be retrieved [11]. This symbolic model is a model of real neural behavioural that has a measurable psychological manifestation in neurons. As a cell assembly is the neural implementation of a psychological concept, the activity in a cell assembly should be correlated to the activation of a psychological concept. The simplest measurement is number

of neurons firing, so if the number of neurons in a cell assembly firing is larger at one time than another, then the corresponding concept is more active.

Cell assemblies are interesting because they can be used to simulate aspects of the functioning of neurological structures, thus there is an urge to make such models biologically plausible. They are also interesting as a novel form of computation; considerable effort is being applied to manufacture processor chips based on cell assemblies as they are parallel processing units which could revolutionise computation if the principles of how to use them to solve real world problems can be elucidated (eg: the EU FACETS project, [12]). Accordingly some cell assemblies stray from strict biological plausibility in pursuit of these criteria.

This paper looks at a number of ways to affect cell assembly decay. The CANT model is a cell assembly model that has variously been used for information retrieval [9], apply rules [8], learning hierarchical categories [7], and as a video game agent among other applications. Currently in the CANT model cell assembly decay is enforced by making cell assemblies inhibit each other or by ending the current training run where all cell activation is reset.

II. The CANT model

The CANT simulator is based on fatiguing Leaky Integrate and Fire (fLIF) neurons. Neurons collect activation from other neurons via synaptic connections. If the neuron does not fire, some of that activation leaks away. Equation 1 describes the activation of a neuron i at time t if it does not fire at time $t-1$.

$$h_{i_t} = \frac{h_{i_{t-1}}}{d} + \sum_{j \in V_i} w_{ji}, 1 < d \tag{1}$$

The amount of leak is given by d the leak parameter. V_i is the set of all neurons that have connections to i and fire at time $t-1$. The weight, or synaptic strength, of the connection from neuron j to neuron i is given by w_{ji}.

The model is based on discrete time steps. This allows the whole system to be updated simultaneously. It can be argued that each time step is roughly equivalent to 10 ms. of simulated time. This enables the system to ignore refractory periods and synaptic delay as these are all within the 10 ms.

Neurons also fatigue so that the more closely adjacent steps at which they fire, the more difficult it becomes for them to fire. While a simplification of the biological system, where fatigue is associated with ion transport, it is modelled by increasing an activation threshold θ if a neuron fires as described by Equation 2.

$$\theta_t = \theta_{t-1} + F_c \tag{2}$$

In Equation 2 the threshold θ at time *t* is set to the threshold at time *t-1* + the fatigue constant F_c. If the neuron does not fire, the threshold is reduced toward the base resting level as in Equation 3.

$$\theta_t = \theta_{t-1} - F_r \tag{3}$$

A. A naive short term memory model.

In order to examine cell assembly temporal dynamics in CANT a naive model was constructed to implement short term memory for numbers. An individual memory could store one number between 0 and 9. Each memory was implemented as a set of ten cell assemblies, of two hundred neurons each, that were connected together. Neurons were inhibitory or excitatory with a ratio of 1 to 4. Inhibitory connections were random. Excitatory neurons have local connections and one long distance axon with several synapses. So a neuron connects to nearby neurons and to neurons in one other area of the net.

The cell assemblies in a memory block were individually trained and learnt to store a number according to the learning rules. They used the standard correlatory learning rule (see [7]). When neurons co-fired the weight was adjusted by equation 4. The weight was reduced according to equation 5 when the pre-synaptic neuron fired and the post-synaptic neuron did not.

$$\Delta^+ w_{ij} = (1 - w_{ij}) * R \tag{4}$$

$$\Delta^- w_{ij} = (w_{ij}) * -R \tag{5}$$

Synaptic weight roughly reflects the likelihood that the post-synaptic neuron fired when the pre-synaptic neuron fired.

The regime for training a memory block was that a high connectivity was specified in order to facilitate quick formation of cell assemblies, and a cell assembly had to form at each number location in the memory block following fifty cycles of activation during which the training stimulus was presented for the first twenty five cycles. If a complete set of cell assemblies did not form then the memory was retrained using a different random number seed until a complete set was formed. When cell assemblies formed for each number the amount of neurons involved in each assembly varied due to the randomisation factors in the construction of cell assemblies and due to competition between assemblies within a memory block during learning. Ten unconnected blocks of memories were constructed so that up to ten numbers between 0 - 9 could be stored. For subsequent storage of numbers learning was turned off and F_c = 0.5, F_r = 1.0, d = 1.1 and θ =4.0. Stimuli (one frame of a hundred and fifty randomly assigned neurons) representing ten random numbers were presented at fifty step intervals, and numbers were stored in the memory blocks using a matrix addressing scheme which meant that successive numbers were stored in successive memory blocks.

In order to examine the firing behaviour of cell assemblies their activity is plotted as a function of time. Given space constraints only the most active (CA1) and least active (CA2) of the ten randomly activated assemblies are shown in figure 2 to give an idea of range. The reverberatory activity of these two assemblies can be clearly seen, as can the difference in magnitude of activity between the them. A rough average (obtained by discarding the first five steps and averaging over the next five hundred) shows CA1 has 91.7 neurons firing per

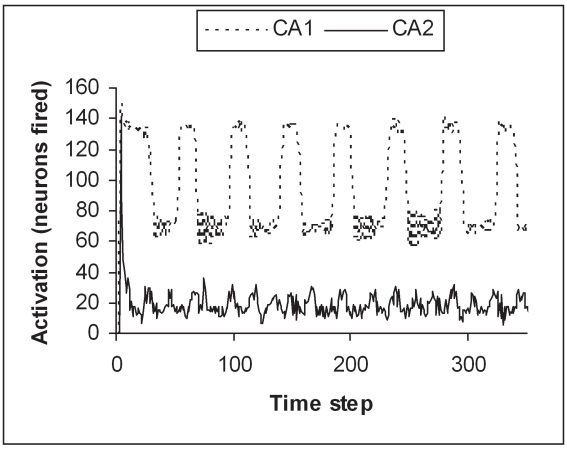

Fig. 2. Activity of two cell assemblies over time.

time frame, while CA2 has only 18.2. Though the numbers of neurons firing in each assembly is quite different they are both well formed and will continue to fire indefinitely in the absence of inhibition from any source.

III. An implementation of Neural Firing Decay

When cell assemblies exist in isolation, the typical behaviour of cell assemblies in CANT models is that neurons either stop firing very quickly or they continue to fire indefinitely. In the first case, neurons in the cell assembly fire; this leads to a cascade of activation (ignition) with all or most of the neurons in the cell assembly firing in each cycle. Fatigue accumulates rapidly in all of the neurons, and eventually some of the neurons do not fire. While these neurons lose some fatigue, they also stop sending activity to other neurons in the cell assembly. This leads to a crash in activity and all of the neurons stop firing.

In the second case, fatigue and activation are more finely balanced. For instance, a smaller percentage of neurons might fire in each cycle after ignition. Fatigue increases in the neurons, but it may increase more slowly as a given neuron does not fire in each cycle. When fatigue causes some neurons to stop firing, they have time to recover. This leads to the case where overall fatigue stays constant as does overall activation. The cell assemblies from the prior section did this.

It is however possible to build a topology where the number of neurons firing decreases in a linear way. Figure 3 gives an example of a system. In this case it consists of 12 neurons $I_1 ... I_6$ and $J_1 ... J_6$. The connections are reciprocal and are weighted from 10 to 5.

The set would be activated by firing the top half. It would then proceed to oscillate between the top and bottom half with six neurons firing in each cycle. If F_r were less than F_c, overall fatigue in the neurons would gradually build up. At some point, overall fatigue in all of the I neurons would be greater than 1 and the I_6 and J_6 neurons would stop firing. As the fatigue in the remaining neurons grows they will cease firing two by two and in a linear fashion.

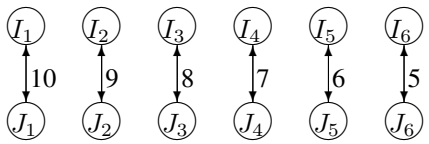

Fig. 3. A Set of Neurons that Oscillates and Decays.

A simulation with 50 pairs of neurons was run. $F_c = 1.2$, $F_r = 1.0$, $\theta = 4.0$, and the connection weights range from 5 to 55. One pair drops out every 10 cycles, with the last pair ceasing at cycle 500. Reducing F_c to 1.1 doubles the period of activation.

A second topology with $F_r > F_c$ has also been developed. The topology is a bit more sophisticated and instead of pairs of neurons, it is based on sets of eight neurons shown in figure 4. The eight neurons are broken up into two groups of four I, J, K, and L. The main connections are shown in both groups in the figure. A neurons sends excitation to its successor of $2.5 + x$ and in this case I succeeds L. The other connections are only shown in the figure on the left side, and are weighted .8. Additionally there is a connection of weight .8 from each neuron to its double in the opposite group (e.g. I_1 to I_2). Initially six of the eight are ignited, typically these are the I, J and K neurons. In the next cycle, the first pair, I, does not get enough energy to fire, but the other three pairs do. So there is an oscillation with six neurons on in each cycle. In the simulation $\theta = 4$, $F_c = 0.01$, $F_r = 0.011$ and $d = 1.5$. The remaining factor is x in the primary connections. The simulation has 20 sets of eight neurons and x increases through each set from 0 to .95 in steps of .05. This means that the first pair fatigues out first, and subsequent pairs fatigue out linearly there after.

A summary of the behaviour of both systems is shown in figure 5. The behaviour shows that both indeed decay linearly over time.

While these two restricted topologies account for linear memory decay, they seem a poor basis for a memory decay model. The topologies are highly restricted and it is hard to imagine how such topologies could be learnt. In the next section, modification of the model leads to more biologically plausible solutions to the memory decay problem. Moreover these system should work over a much wider range of topologies so that the topologies can be used to encode different information.

IV. Decay from long term fatigue

Possibly the simplest way to implement a long term reduction in neural firing in a cell assembly is to implement a

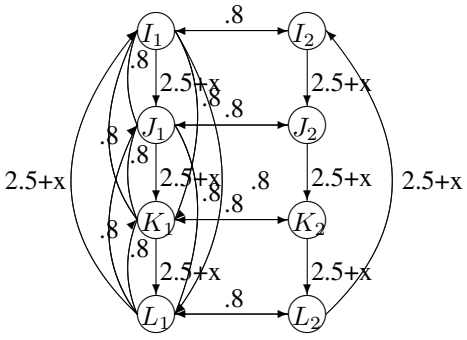

Fig. 4. A Second Set of Neurons that Oscillates and Decays.

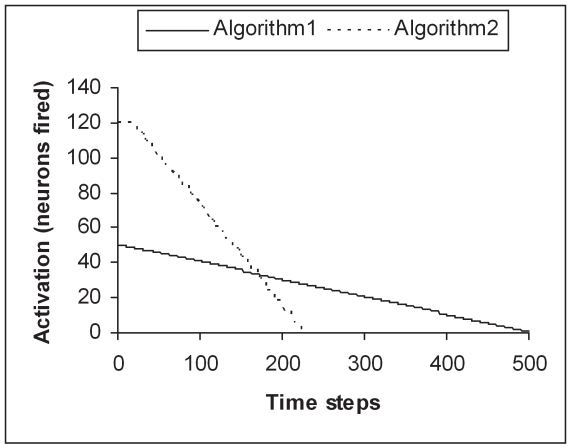

Fig. 5. Neural Firing over time in two restricted topologies

form of long term fatigue. Such fatigue should reduce activity by raising the activation threshold as is done by the existing fatigue parameter. It should be noted that the TRACE model [10] implements this sort of double fatigue in a mathematical model of cell assembly behaviour. Here, this mathematical model is translated into a neural model where assemblies emerge. Long term fatigue was implemented so that equations 2 and 3 are replaced by the following equations. LF_r denotes the long term fatigue recovery rate and LF_c the long term fatigue constant.

$$LF_{i_t} = LF_{i_t-1} + LF_c \quad (6)$$
$$LF_{i_t} = LF_{i_t-1} - LF_r \quad (7)$$

Short term fatigue SF can then be redefined in a similar manner.

$$SF_{i_t} = SF_{i_t-1} + SF_c \quad (8)$$
$$SF_{i_t} = SF_{i_t-1} - SF_r \quad (9)$$

The activation threshold θ is then modulated by the values of these two fatigue factors.

$$\theta_t = \theta_{t-1} + SF + LF \quad (10)$$

Figure 6 shows the effect on activation of implementing long term fatigue. The same cell assemblies as used in figure 1 were

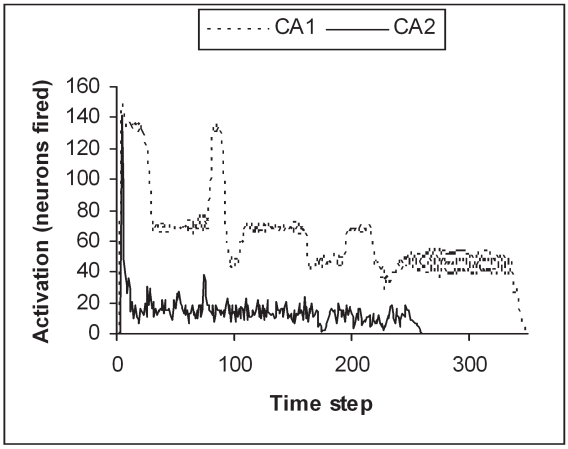

Fig. 6. Activity of two cell assemblies over time with long term fatigue.

activated using the same random seed as for figure 1 but using the long term fatigue model. LF_c was set at 0.004 and LF_r at 0.0002 (and the short fatigue parameters as before). The activity of both cell assemblies reduces and now extinguish as a function of time in a manner reminiscent of figure 1. It is interesting to note that CA2 dies off around time step two hundred and fifty and CA1 at timestep three hundred and fifty which is relatively close given they have such different levels of overall activation.

A. Effect of long term fatigue on response to an ambiguous stimulus

In the previous section it was demonstrated that long term fatigue can extinguish cell assembly activity. The next step is to investigate the temporal dynamics over a longer time scale. One aspect of this is how competing fatiguing assemblies interact, which should be demonstrated by the presentation of an ambiguous stimulus. The Necker cube is probably the most famous example of an ambiguous or bistable stimulus. In this case a line drawing that can be interpreted as a 3D wireframe cube seen under orthographic projection can be perceived as being in two possible orientations. Subjects usually report that they see one orientation only at a time and that from time to time the percept flips from one orientation to the other. While there are many explanations and influences on the perception of bistable stimuli, fatigue is one explanation - that two percepts are competing and that whichever is currently dominant will eventually give way to the other due to fatigue. In a much more simple analogue of the bistable stimulus two cell assemblies were formed (two hundred neurons each) by presentation and training on two stimuli, an upper rectangular block and a non overlapping lower rectangular block. Subsequently learning is turned off and an ambiguous stimulus (a number '1' which overlapped the positions of the upper and lower blocks roughly equally was presented continuously and the cell assembly that dominated was periodically recorded as dominance oscillated between the two assemblies. The parameters were LF_c = 0.0005, LF_r = 0.0001, SFc = 0.5,

SFr =1.0, d =1.1 and θ = 4.0.

The result was that the system produced osccillations in the expected manner. One cell assembly was slightly stronger than the other as it had recruited more neurons during training and thus dominated slightly more in the long term, but because it fatigued the other cell assembly won some of the time. If a cell assembly being active is arbitrarily defined as having any cells firing at particular intervals (in this case at fifty cycle intervals from cycle twenty five), then over one thousand runs the top cell assembly was dominant for five hundred and forty five cycles, the bottom for four hundred and forty three, and in two cases recording was made on transition between the two states. For these parameters the system transitioned from one state to the other eleven times.

V. Decay from activation leak.

In the long term fatigue condition decay comes about because neurons firing thresholds are raised. An alternative approach to decay would be to model the reduction of activation via leak. In the CANT model the leak of activation is modelled by the parameter d in equation 2. d is a constant and its value must be greater than 1. For a given cell assembly there is a value for d which so quickly undermines activation that the cell assembly is extinguished. We can model leak as a factor that increases with neuronal firing, in which case leak can be uses as a means for model extinction. We define parameters for leak similar to those for fatigue by having a leak constant d_c added to the base rate of leak when a neuron fires, as in equation 11, and a leak recovery parameter d_r which allows the leak rate to return to the base rate when the neuron does not fire, as in equation 12.

$$di_t = d_{i_t-1} + d_c \tag{11}$$

$$d_{i_t} = d_{i_t-1} - d_r \tag{12}$$

This scheme was then implemented with the values d_c = 0.001, d_r =0.0005, $F_c = 0.5$, $F_r = 1.0$, $\theta = 4.0$, and the decay base rate was set to 1.1. The model was then run using the same random seed on the short term memory model described in section IV, on the same two cell assemblies. The results are shown in figure 7. The graph shows that both assemblies are extinguished after some time although now it is CA2 which persists for longer. To get a rough estimate of the variability in extinction rates for the long term fatigue and leak conditions twenty runs each with different random numbers seeds were made for both CA1 and CA2. The results are shown in table I. The fact that the standard deviations are relatively small suggests that these methods of extinction may be fairly reliable. It is noticeable that the effect of size of cell assembly may have opposite effects under the two conditions as CA1 tends to last longer than CA2 on the long term fatigue condition but shorter on the leak condition. This could bear further investigation if either of these techniques are found to be worth using.

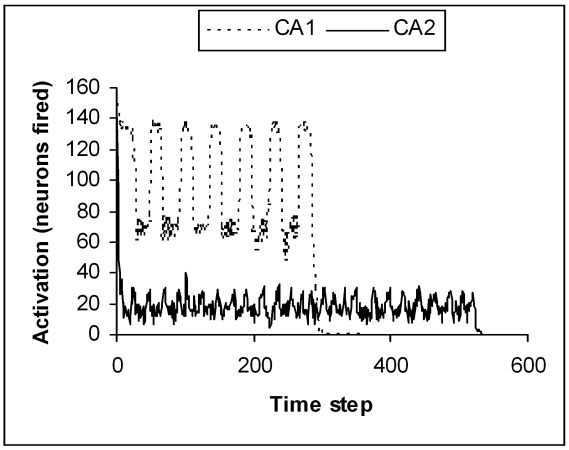

Fig. 7. Activity of two cell assemblies over time with activation leak.

TABLE I
MEAN TIME STEPS TO EXTINCTION WITH STANDARD DEVIATIONS FOR 20 RANDOM RUNS FOR CA1 AND CA2 WITH LONG TERM FATIGUE AND ACTIVATION LEAK.

	Long term fatigue		Activation leak	
	mean	std. dev.	mean	std. dev.
CA1	297.4	49.3	253.8	44.4
CA2	456.5	82.3	711.2	197.0

A. Effect of activation leak on response to an ambiguous stimulus

To investigate one aspect of longer term dynamics of activation leak the ambiguous stimulus experiment run for long term fatigue was run under the condition of activation leak using the parameters given in the previous section. Again consecutive domination of one cell assembly over the other was seen, as each assembly's activation successively leaked giving way, allowing the other assembly to prevail. Over a thousand cycles with the given parameters the system transitioned from one assembly dominating to the other thirty eight times; for four hundred and fifty cycles the bottom CA dominated, five hundred forty three times the top assembly dominated, while seven times the system was registered as in transition. Therefore as predicted a similar temporal dynamics was found on this task as found for the long term fatigue example.

VI. The effect of refreshing the stimulus.

Both the long term fatigue and increasing leak approaches to producing decay in cell assemblies seem quite effective and stable from a preliminary analysis, and may be more biologically plausible than the methods presented in section III. However one test they may fail is that of being able to be refreshed over long periods. Some behaviour of memory need the capability to have activity refreshed by re-presentation of the stimulus. An example of this is the ability to remember a phone number that has just been heard for longer than

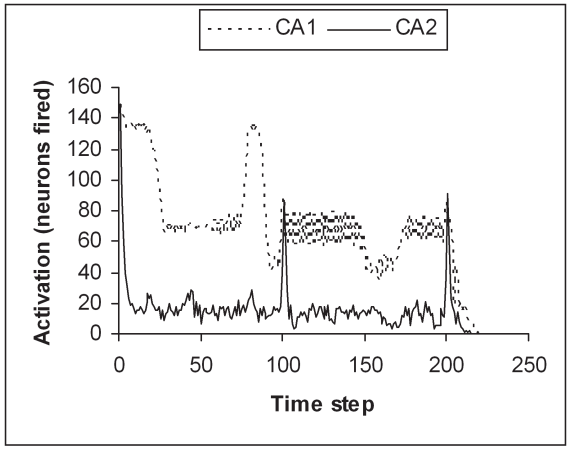

Fig. 8. Activity of cell assemblies over time which are fatiguing and being refreshed every hundred cycles.

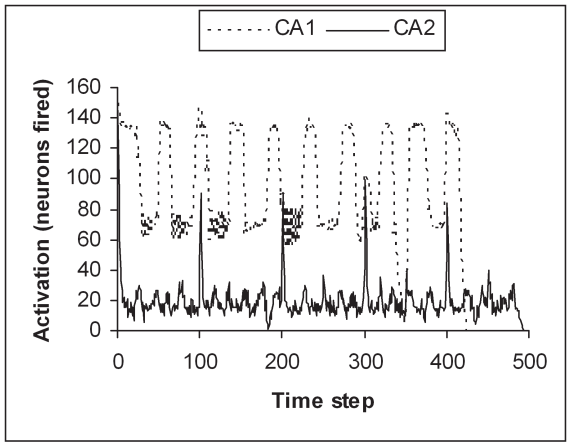

Fig. 9. Activity of cell assemblies over time which are leaking activation and being refreshed every hundred cycles.

a few seconds, by mental rehearsal. Accordingly, the two example cell assemblies were re-tested with the stimulus being presented on the first and then once every hundred frames until extinction. The results are shown in figures 8 and 9 respectively. In the fatigue condition activation is shortened for both assemblies, in the leak condition activation is lengthened slightly, but in all cases extinction is inevitable. These results suggest that for the parameter range tested, on their own neither of these two methods can reproduce the desired effects of refresh. It remains to be seen whether such methods may work in conjunction with other mechanisms to produce the required behaviour.

VII. Conclusion

This paper has investigated two classes of methods for producing autonomous decay in cell assemblies. In the first class two cell assemblies whose activation decays in a highly precise linear manner, due to the nature and strengths of their connections, were examined. Whilst it is interesting to note the precision with which this decay is produced from a connectionist programming point of view, it is hard to imagine how such circuits could be learnt in a biological system.

The second approach, which may be more biologically plausible, looked at manipulation of variables within the CANT cell assembly model. Both activation decay and long term fatigue produced reasonable passive decay in activated assemblies. They both produced reasonable behaviour over time in their response to ambiguous stimulus where assemblies successively rise to dominance despite differences in initial strength, due to the effects of decay. However neither of these two approaches, as investigated here, can account for a crucial aspect of temporal dynamics: the ability to maintain activation on stimulus refresh. In the ACT-R model for example [11], repeated activations lead to longer activations. Future work will address this shortcoming in the current model.

Acknowledgment

This work is supported by EPSRC grant EP/D059720 - Natural Language Parsing using Cell Assemblies: Computational Linguistics with Attractor Nets.

References

[1] M. Abeles, H. Bergman, E. Margalit, and E. Vaadia. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, 70(4):1629–38, 1993.

[2] A. Baddeley. *Working memory*. Oxford University Press, Oxford, 1986.

[3] N. Burgess and G. Hitch. Memory for serial order: A network model of the phonological loop as a language learning device. *Psychological Review*, 106:551–581, 1999.

[4] T. Preece G. Brown and C. Hulme. Oscillator based memory for serial order. *Psychological Review*, 107:127–181, 2000.

[5] S. Grossberg. Towards a unified theory of neocortex: laminar cortical circuits for vision and cognition. Technical Report CAS/CNS TR-2006-008, Department of Cognitive and Neural Systems, Boston University, USA, 2006.

[6] D. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York, 1949.

[7] C. Huyck. Creating hierarchical categories using cell assemblies. *Connection Science*, 19:1:1–24, 2007.

[8] C. Huyck and R. Belavkin. Counting with neurons: Rule application with nets of fatiguing leaky integrate and fire neurons. In *Proceedings of the Seventh International Conference on Cognitive Modelling*, pages 142–147, 2006.

[9] C. Huyck and V. Orengo. Information retrieval and categorisation using a cell assembly network. *Neural Computing and Applications*, 14:282–289, 2005.

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

[11] R. Lewis and S. Vasishth. An activation-based model of sentence processing as skilled memory retrieval. *Cognitive Science*, 29(3):375–419, 2005.

[12] O. Rochel S. Chemla P. Kornprobst T. Viville A. Daouzli S. Saighi C. Lopez and S. Renaud. A first step towards in silico neuronal implementation of early-vision map. In *SSD'07: Fourth International Multi-Conference on Systems, Signals and Devices*, 2007. ISBN: 978-9973-959-06-5.

[13] P. Milner. The cell assembly: Mark ii. *Psychological Review*, 64:4, 1957.

[14] M. Page and D. Norris. The primacy model: a new model of serial recall. *Psychological Review*, 10:761–781, 1998.

[15] F. Pulvermuller. Words in the brain's language. *Behavioral and Brain Sciences*, 22:253–336, 1999.