

Effects of compensation, connectivity and tau in a computational model of Alzheimer's Disease

Rowan, Mark

DOI:

[10.1109/IJCNN.2011.6033268](https://doi.org/10.1109/IJCNN.2011.6033268)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Rowan, M 2011, 'Effects of compensation, connectivity and tau in a computational model of Alzheimer's Disease', 2011 International Joint Conference on Neural Networks (IJCNN 2011 - San Jose), San Jose, CA, USA, 31/07/11 - 5/08/11 pp. 543-550. <https://doi.org/10.1109/IJCNN.2011.6033268>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Institute of Electrical and Electronics Engineers

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Effects of Compensation, Connectivity and Tau in a Computational Model of Alzheimer’s Disease.

Mark Rowan, `m.s.rowan@cs.bham.ac.uk`

School of Computer Science, University of Birmingham, B15 2TT, UK

Abstract—This work updates an existing, simplistic computational model of Alzheimer’s Disease (AD) to investigate the behaviour of synaptic compensatory mechanisms in neural networks with small-world connectivity, and varying methods of calculating compensation. It additionally introduces a method for simulating tau neurofibrillary pathology, resulting in a more dramatic damage profile. Small-world connectivity is shown to have contrasting effects on capacity, retrieval time, and robustness to damage, whilst the use of more easily-obtained remote memories rather than recent memories for synaptic compensation is found to lead to rapid network damage.

I. INTRODUCTION

Alzheimer’s disease (AD) is a specific form of dementia, characterised biologically by neurofibrillary tau protein tangles and beta-amyloid ($A\beta$) protein plaques [28], and symptomatically by a progressive decline in memory capabilities. In particular, recent memories are the first to be lost whilst distant memories are retained, but as the disease progresses this is followed by gradual total loss of recall, a corresponding loss of personality, motor control, and other bodily functions, and finally death [8].

Computational modelling of neurological disorders such as AD is an established tool [1] but existing models of AD such as [11], [23] can now be improved in line with better understanding of the disease. One such model, by Rupp and Reggia (1995) [23], showed how simple lesions in a single-layer associative network trained in an *activity-dependent* Hebbian manner leads to loss of memory, and the addition of a local compensation factor causes the pattern of functional damage to mirror more closely that found in AD whereby recently-stored memories are lost before historical memories. Later work showed how the compensation factor can be made biologically plausible by depending only on the post-synaptic potential of the remaining neurons after lesioning [13].

This model remains widely cited [5], [25] even though it could be made capable of better approximation of the lesions representing AD pathology; currently it only either deletes neurons and synapses at random or deletes neurons within a specified radius on a 2-D grid [24]. Today we know much more about connectivity strategies in the brain such as small-world networks [30], [4] as well as the biological processes underpinning AD, such as neurofibrillary tau pathology.

In this paper, methods are presented for enhancing the Rupp and Reggia model with up-to-date techniques which may be more representative of the underlying biology. This work is intended to examine differences in behaviour which may occur when considering connectivity strategies, specific

details of compensatory techniques and lesioning in accordance with specific pathologies, with the aim of leading to development of more accurate representations of a range of pathological processes underlying AD such as those involving tau, beta-amyloid, and N-amyloid precursor protein [20], in more complex network models such as LEABRA [21], spiking neural networks [9], and reservoir networks [17].

The remainder of this paper is organised as follows: Section II describes the Rupp and Reggia model in greater detail and the updates made to it in this work, section III presents the results of experiments characterising the network’s behaviour with these new enhancements, and section IV deals with concluding remarks and outlines future directions in which this research could be taken.

II. MODEL DESCRIPTION

A. Learning rule

Rupp and Reggia showed how a variant of an attractor network model proposed by Tsodyks and Feigl’Man (the T-F model) [29] is capable of storing patterns in a biologically-plausible Hebbian *activity-dependent* manner. This is achieved using a repetitive-learning process whereby each pattern to be stored “must be presented to the network several times before it becomes engraved on the synaptic matrix with sufficient strength, and is not simply enforced on the network in a ‘one-shot’ learning process” [23]. An updated version of the model [24] added Gaussian partial-connection of the network rather than full connectivity.

$$W_{ij}(t) = W_{ij}(t-1) + \frac{\gamma}{N}(S_i - p)(S_j - p) \quad (1)$$

The network learns patterns through a process of activity-dependent learning according to the update rule in equation 1. A set of *external inputs* delivers activation greater than the neural threshold to each unit of the network according to the pattern to be learned. W is the weight matrix of undirected connections between neurons i and j , γ is a constant determining the magnitude of activity-dependent changes, N is the number of neurons in the network, $K \leq N$ is the number of other units to which each unit is connected, S refers to the neuronal state $\{0, 1\}$, and p is the coding rate denoting the proportion of 1s compared to 0s in the stored memory patterns ($p \ll 1$ as cortical networks are found to have low coding rates [2]).

The activity-dependent learning rule for pattern storage is based on the Hebbian principle but introduces the requirement for each given pair of units to remain in the same state for a certain number of update iterations (the suggested value is 5) before the synaptic weight between them is updated, and requires each pattern to be presented several times in turn to the network before it is completely stored. Thus the learning algorithm attempts to mitigate the effects of the Hebb rule's ability to globally alter synaptic weights in a biologically-unrealistic way and circumvents its method of storing each pattern in a 'one-shot' process which is susceptible to the presence of errors or noise. By presenting each pattern several times to the network, any noise present in the inputs is reduced and the synaptic matrix is gradually constructed rather than being enforced in a single process by the learning rule.

B. Performance evaluation

Patterns are recalled using a noisy version of the complete pattern applied to the network via the same set of external inputs used for learning with activation less than the neural firing threshold. A measure of the recall performance the network for a given pattern ξ^μ , termed the *overlap* between the resulting network state and the pattern, has the effect of counting the correctly-firing units whilst also penalising with a lower weighting those units which fire erroneously (equation 2) [29]:

$$m^\mu(t) = \frac{1}{p(1-p)N} \sum_{i=1}^N (\xi_i^\mu - p) S_i(t) \quad (2)$$

C. Synaptic compensation

In the work by Rupp and Reggia the network model was lesioned by deleting synapses or neurons at random and implementing a process of *variable synaptic compensation*, where "the magnitude of the remaining synapses is uniformly strengthened in a manner that partially compensates for the decrease in the neuron's input field" [23] by multiplying the weights of the remaining synaptic connections by a globally-determined (i.e. depending on knowledge of the overall fraction of deletion) local compensation factor.

Rupp and Reggia examined the overall degradation in recall performance and the pattern of relative sparing of older memories compared to recently stored patterns (as observed in AD patients [15]) as the network was progressively lesioned, and concluded that synaptic deletion and compensation in this model can be demonstrated to reveal similar symptoms to the cognitive decline observed in AD.

However a global synaptic compensation strategy is biologically implausible as each neuron must somehow be aware of the global deletion rate both for itself, and for other neurons around it. Horn et al. [13] therefore introduce a neuronal-level compensatory mechanism which causes each neuron to adjust its output based only on changes in the neuron's average post-synaptic potential (or summed input), and which does not require the explicit knowledge of either global or local levels of synaptic deletion.

At any given moment, each neuron has an estimate \hat{w}_i of its total connectivity compared to the starting value ($w_i = 1$). It can compensate for this reduced connectivity by multiplying the remaining incoming synapses (essentially, lowering its firing threshold) by a value c_i . This is achieved via repetition of the following steps:

- In the pre-morbid state (i.e. before each iteration of lesioning) a set of random noise patterns ($p \ll 1$) is presented to the network and it is allowed to fall into a stable state. Each neuron then obtains its resulting input field measurement, the expected value of which is denoted $\langle h_i^2 \rangle$.
- Horn et al. state that $\langle h_i^2(\hat{w}_i) \rangle = c_i^2 \hat{w}_i \langle h_i^2(w_i = 1) \rangle$. Given an assumption that $c_i^2 \hat{w}_i = 1$ (i.e. the network is currently correctly compensating for any value of $w < 1$), the neuron's average "noise-state" input field value, the expected value of which is denoted $\langle R_i^2 \rangle$, is therefore equivalent to $\langle h_i^2 \rangle$.
- The same process is repeated using a set of already-stored patterns rather than random noise patterns. Each neuron then obtains its resulting average "signal-state" input field strength, the expected value of which is denoted $\langle S_i^2 \rangle$ (Horn et al. speculate that this process could occur biologically during dreaming). As with the earlier noise term, $\langle S_i^2 \rangle \equiv \langle h_i^2 \rangle$ as $c_i^2 \hat{w}_i = 1$.
- The network is lesioned in some way unknown to the individual neurons (e.g. by deleting synapses).
- Now, in order to estimate the new value of \hat{w}'_i in the post-morbid state, and thus to compute a new value for c'_i , a further set of already-stored patterns is presented to the network and the network allowed to converge once more to a stable state. A new *post-morbid* value for each neuron's input field $\langle h_i^2 \rangle$ is obtained.
- Horn et al. separate this $\langle h_i^2 \rangle$ into signal and noise terms: $\langle h_i^2(w_i) \rangle = c_i^2 \hat{w}_i^2 \langle S_i^2 \rangle + c_i^2 \hat{w}_i \langle R_i^2 \rangle$. The noise term is already known from the earlier steps, and is subtracted from the post-morbid input field value. It is thus possible to calculate \hat{w}'_i using equation 3, and then to derive a new value for the compensation c'_i :

$$\hat{w}'_i = \frac{\langle h_i^2 \rangle - \langle R_i^2 \rangle}{c_i^2 \langle S_i^2 \rangle} \quad (3)$$

D. Unanswered questions

Whilst experimental support exists for the predicted compensatory strengthening of synapses in AD [25], one unexplained result drawn from this model is that significant neuronal deletion (around 50%) in the model is required before memory function is seriously impaired. This rate of deletion is much larger than the rate observed clinically in the latest stages of Alzheimer's disease (between 10% and 30% reduction of volume in the hippocampal regions in severe cases of AD [18] and certainly far more than the 10% general cerebral atrophy reported at initial diagnosis of the disease [22]), implying that there must be other factors additionally affecting cognitive decline.

A further limitation of the model is that lesioning is performed only by deleting a number of randomly-selected neurons or connections at each step, which does not necessarily represent the subtleties of the underlying pathology. The authors present a method of applying lesions in a localised spatial manner by deleting all of the neurons and/or connections within a circle or rectangle of a given area [24], but this does not incorporate any of the known neurodegenerative mechanisms such as tau or amyloid pathology.

During the synaptic compensation process previous studies have not examined the differences in performance when using recent versus remote memories to calculate the signal term.

Finally, the method in which the network is interconnected (either fully, or using an arbitrary number of connections per neuron in a localised Gaussian manner) is again simplistic and does not represent biologically realistic connection strategies such as small-world networks [30] or neural Darwinism (pruning of weaker synapses during development) [12].

E. Implementing different connectivity strategies

Biological neural networks such as those found in the hippocampus are generally sparsely connected [16], [4]. It has been shown that in Alzheimer's disease, small-world clustering (as measured by the clustering coefficient) is significantly reduced at a global level, resulting in large changes to the local organisation of the network [27].

Connection dilution mechanisms for associative networks include connecting each unit over a flat random distribution, wiring each unit in a spatial manner to those immediately surrounding it with Gaussian probability, and using a randomised small-world network connection strategy [30].

Small-world networks in this model are constructed in the form prescribed by Watts and Strogatz [30] by firstly connecting each neuron to its closest K neighbours. Then, according to a probability of re-wiring $p(\text{rewire})$, the connections between each unit and its two immediate neighbours are randomly assigned to other units in the network. Once each unit in the network has been considered, the neighbours two places away from each unit are then considered, and then those three places away, until each connection in the network has finally been randomly re-wired or left in place.

It has been shown that a process of diluting the synaptic weight matrix of a Hopfield network such that it is no longer fully-connected still causes it to behave in much the same manner as a fully-connected network when stored state vectors are generally of low activity [6].

F. Implementing tau lesioning

1) *Medical background:* The tau hypothesis refers to the neurodegenerative effects of a modified (or *hyperphosphorylated*) form of the tau protein which aggregates with other fibres of tau and eventually forms the neurofibrillary tangles (NFTs) inside neurons which are prevalent in brains with AD. It proposes that the cognitive decline in AD is due primarily to loss of synapses and neurons (via a toxic form of the modified tau), and the subsequent loss of connectivity experienced [26].

The normal function of tau is disrupted in AD: "tau is essential for establishing neuronal cell polarity and axonal outgrowth during development and for maintaining axonal morphology and axonal transport [of neurotransmitter-containing vesicles along the axon] in mature cells" [7], [14] and in both constructing and stabilising microtubules which, in developing neurons, are important for establishing neuronal cell polarity and outgrowth, and in adult neurons are essential for proper structure, function and viability [7]. Instead of binding to the microtubules, tau in AD becomes sequestered into NFTs within the neurons [3] and as the level of normal tau in the brain is reduced the microtubules disintegrate, causing further neuronal dysfunction. The existence of NFTs could also present a toxic gain-of-function by physically obstructing the transport of vesicles within the neuron (leading to cognitive impairment) and also by further sequestration of normal tau into the modified form as part of a cascade of neurodegeneration [3].

Although the amyloid hypothesis provides a possibly more widely encompassing view of AD, it has a number of significant unexplained problems, not least that "the number of amyloid deposits in the brain does not correlate well with the degree of cognitive impairment" [10], and so the tau hypothesis and its relationship to the amyloid hypothesis remain an important subject for further research.

2) *Computational implementation:* Whilst it is possible to delete either neurons or the connections between them as shown in earlier studies [23], these processes can be considered essentially the same: if all the incoming synapses of a particular neuron are deleted, the neuron is no longer able either to receive activation from surrounding neurons, nor is it able to have any excitatory effect on its neighbours. This neuron might just as easily be considered to have been deleted, as it is effectively removed from the network completely.

Hopfield-type networks (including the T-F network) also suffer from the inherent problem that the output layer is essentially the only layer of the network. That is, whenever lesions are applied to the network by deleting neurons, this necessarily results in the inability of the network to completely recover a cued pattern regardless of the effects the removal of these neurons may have had on the underlying pathology of the network, as the 'output layer' is now only partially complete and can no longer map with full accuracy to every given pattern, resulting in a perceived decrease in network performance.

So it is suggested that results more representative of the underlying pathology could be obtained by the introduction of a subtle shift from *synaptic loss* to *neuronal atrophy*, whereby whole groups of synapses with a single neuron at their centre are affected in a similar way at the same time, but without fully removing the synapses or any neurons from the resulting output patterns used for evaluating network performance.

With this in mind, neurofibrillary tangles (NFTs) of hyperphosphorylated tau are known to result in direct blocking of axonal transport [3] and collapse of microtubules

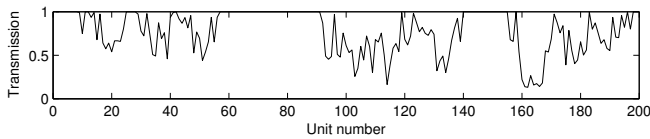


Fig. 1: Characteristic plot of neural damping levels after spatial tau lesioning.

also supporting axonal transport [7]. In order to model more specifically the pathological effects of tau NFTs it is suggested that, rather than simply deleting neurons or connections at random or in areas of a certain radius [23], [24], the output of selected neurons could be partially muted to simulate the effects of axonal blocking by NFTs.

To simulate the sequestering of hyperphosphorylated tau and the subsequent cascading spread of damage, neighbouring neurons could also be muted by a slightly smaller amount, with new tau lesion centres subsequently formed near to existing lesions, and the resulting distributed damage occurring in less of a severe ‘binary’ manner as with random deletion of synapses or neurons.

In computational terms, lesioning can be performed in steps of size z . The locations of the centres of the first set of lesions (the tau *seeds*) are chosen at flat random from across all neurons and are assigned to set D . Subsequent lesioning steps proceed as follows:

- A subset $d \subseteq D$ of size z locations are chosen at flat random. With Gaussian probability centred on each element in d , a new set d' of neighbours is chosen.
- Each neuron’s activation is dampened by multiplying by a value drawn from a Gaussian distribution as a function of the neuron’s proximity to the lesion centre (i.e. $x - \mu$), and with lesion width σ (suggested as 2), such that those neurons closest to the lesion centres in d' are most heavily diluted, and those distant from the elements of d' are relatively unaffected.
- d' is added to D and the process is repeated with the new, larger set of lesion centres D . This results in characteristic lesioning as seen in figure 1.

III. RESULTS

Unless otherwise stated, all experiments were performed in a network with the following parameters: network size $N = 1600$, connections per unit $K = 200$, neural threshold $\theta = 0.048$, noise $T = 0.005$, learning rate $\gamma = 0.025$, external input strength (learning mode) $e_l = 0.065$, external input strength (retrieval mode) $e_r = 0.035$, coding rate $p = 0.1$, deletion step $\Delta d = 0.01$. Results were averaged over 10 runs and the number of patterns stored on each run was 10.

A. Random deletion and local field-dependent compensation

In the first experiment, an attempt was made to replicate the results of Ruppín and Reggia [23] using the improved neural field-dependent compensation rule of Horn et al. [13], in which compensatory mechanisms extended the working life of the network during repeated synaptic deletion. A set of 10 patterns was stored in the network and the average

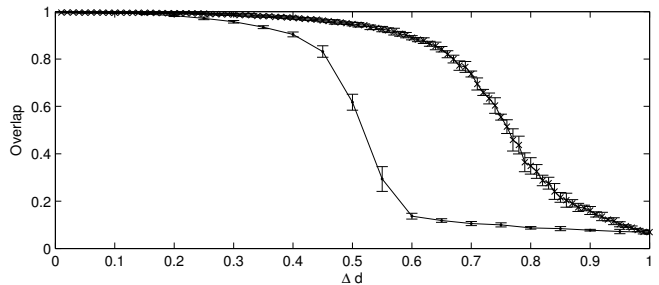


Fig. 2: Performance over synaptic deletion without compensation (leftmost curve) and with (rightmost curve). The deletion step size Δd is larger for the curve without compensation, but this does not affect the overall result as deletion step size-sensitivity is only introduced with compensation.

retrieval success rate (*overlap*) after various levels of deletion was plotted. The results shown in figure 2 are comparable with those achieved by Horn et al. [13] and can be used as a performance baseline for later experiments. The network was connected with a Gaussian connection strategy.

B. Compensation using recent versus remote memories

Ruppín and Reggia observed a gradient of damage by repeating a process of learning a set of patterns then subsequently deleting a proportion of the connections between units [23]. Their results, based on a fixed synaptic compensation strategy, showed a clear decrease in recall performance for patterns learned recently compared with those learned earlier in the process. These results were replicated, and are shown in figure 3.

As the local field-dependent compensation strategy of Horn et al. [13] works by using the retrieval of stored memories for comparing average post-synaptic potentials before and after damage, the choice of memories which should be used for this purpose becomes significant due to the different retrieval success rates of patterns stored early in the lesioning process compared to those stored more recently.

As shown in figure 3a, if only remotely-stored patterns are used during the compensatory process (dotted line) the performance of the network is severely degraded even at a relatively low level of deletion, whilst use of the most recently-stored patterns during compensation is almost indistinguishable in performance from the results when using a random set of patterns drawn from all those previously stored (dot-dash and dashed lines, respectively).

In these cases, a clear gradient of learning has been observed such that patterns stored most remotely are recalled more successfully than those stored more recently (note that this is a distinct phenomenon from serial-position effects in which recency and primacy of items within a list correlate with greater recall, as the network is storing time-separated sets of patterns in between periods of damage rather than a single list of items). At higher, catastrophic levels of deletion (figure 3b), compensation using randomly-selected patterns slightly outperforms compensation using only the most recently-stored patterns, but within the margins of error.

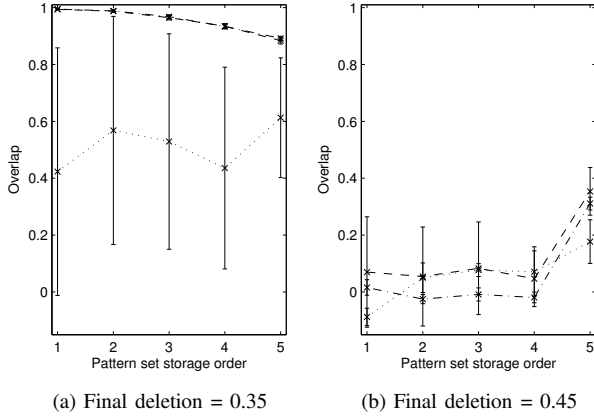


Fig. 3: Performance on separately stored sets of memories. The network was alternately presented with sets of 6 patterns then subjected to a process of deletion with compensation using only the first set of patterns stored (dotted line), only the last set of patterns stored (dot-dash line), or using a random set of 6 patterns drawn from all those previously stored (dashed line). By the final round, the total proportion of deletion was either 0.35 (fig 3a) or 0.45 (fig 3b).

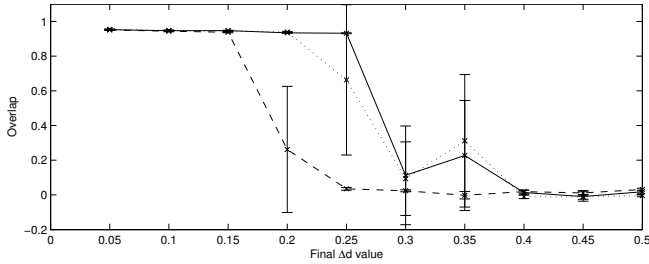


Fig. 4: Average performance on all learned sets (5 sets of 10 patterns, $N = 1200$, $K = 200$) after different final levels of deletion. First-set compensation (dashed line) performance declines significantly earlier than latest-set compensation (dotted line), which itself is marginally less robust than random-set compensation (solid line).

Despite the greater accuracy of their recall within a functioning network, using only remotely-stored patterns during compensation results in much earlier decline of the network performance as deletion progresses (figure 4). Conversely, compensation using the most recently-stored patterns, or sets of patterns drawn at random, results in greater robustness to damage. This could be due to the effect of decreased variance in the patterns used to calculate the signal term during compensation when using only a small, fixed set of patterns, resulting in increased noise during the compensatory process. Any noise arising from using only the first learned set during compensation is multiplied on each compensatory step due to the lower input variance. When using the latest set of stored patterns or a random set at each compensatory step, the variance of the data is increased and this helps to keep noise to a minimum.

This has implications for sufferers of Alzheimer’s disease.

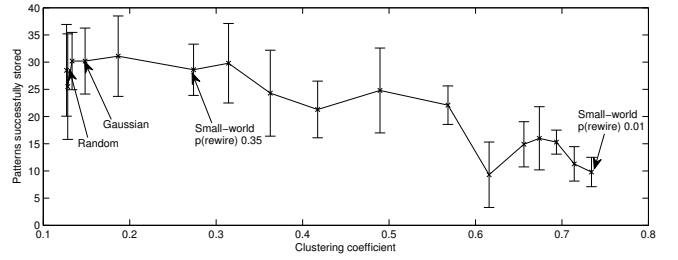


Fig. 5: Capacity in networks with different small-world clustering coefficients. Flat-random, Gaussian and representative small-world networks are indicated individually.

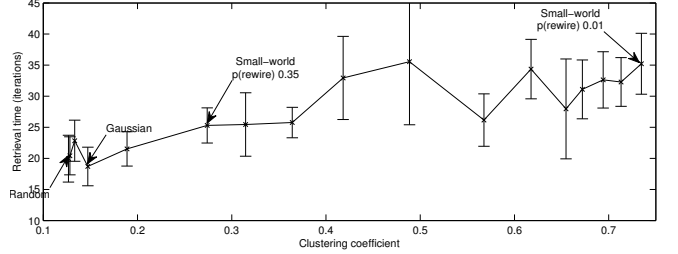


Fig. 6: Retrieval times in networks with different small-world clustering coefficients.

If the network is damaged to such an extent that the gradient observed in figure 3a (the uppermost dashed and dot-dash lines) is evident, but the network is not yet catastrophically damaged, then there may be a greater likelihood that compensatory mechanisms will use remotely-stored patterns compared to recently-stored ones due to their higher recall success. As this experiment has shown, this could actually lead to earlier overall decline of cognitive abilities, and a cycle of correspondingly worse recent memory retrieval performance and thus continued use of remote memories during compensation. Additionally, this finding places the deletion threshold of the model (beyond which all recall is severely affected) closer to the reported 10-30% atrophy levels seen before the symptomatic damage evident in AD.

C. Connection strategies

1) *Effects on network capacity:* Next, the effects on network capacity and robustness to damage of various connection strategies were compared. Firstly, networks were created with $N = 800$ units with connection density $K = 0.125N$. The networks were wired with Gaussian, flat-random, and small-world (with various values for $p(\text{rewire})$) connectivity. Patterns were stored in each network according to equation 1 and retrieved immediately after storage. The average retrieval success rate was plotted against the small-world clustering coefficient of the network’s connection matrix, and when the average overlap measure dropped continuously below 0.8, the network was assumed to have reached its capacity (figure 5).

The results indicate that specific network connectivity generally has little effect on capacity, with the random and Gaussian networks appearing on the same trend as the small-world networks, but it appears that network capacity is significantly reduced in more highly-ordered networks with high clustering

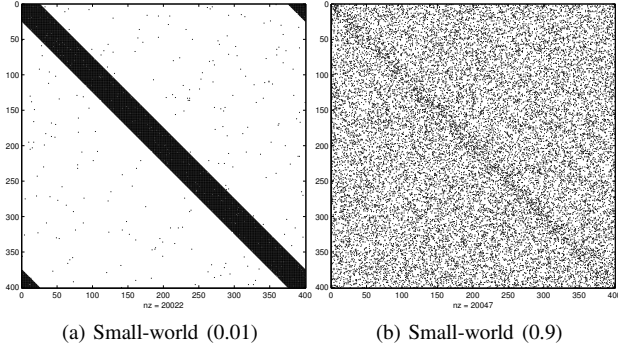


Fig. 7: Connection matrices of networks connected in a small-world manner ($N = 1600$, $K = 200$, $p(\text{rewire}) = 0.01, 0.9$, only every fourth connection plotted for clarity).

coefficients such as small-world networks with low values for $p(\text{rewire})$. Although each network structure contains the same number of connections as the others and should therefore have effectively the same capacity, the difference becomes clear when examining the layout of the connections and the related small-world clustering coefficients. Compare figures 7a and 7b, both of which were constructed using the small-world algorithm. The majority of connections in the network with $p(\text{rewire}) = 0.01$ are located incredibly densely within the local neighbourhood of each neuron, with only a few projections to more distant parts of the network, resulting in a clustering coefficient of 0.73.

This appears to have two effects: the first becomes clear when considering the pattern recall times in figure 6, which shows that in highly-regular networks the number of iterations required for the network to fall into a stable state is higher. This is likely to be due to the lack of distant projections to other parts of the network: activation is ‘slowed-down’ by having to flow through a closely-linked chain of units from one extent of the network to the other, whilst a network with less regularity and more distant projections (as seen in figure 7b) can effectively take short-cuts when activation to distant parts of the network is required. If this activation degrades over time as it traverses the network in small steps, or if there is a limit to the permitted time between cueing and retrieval of a pattern, it is clear to see that these effects could result in greater retrieval failure rates (and thus lower effective capacity) than in a network with less regularity in its connection matrix.

2) *Effects on redundancy and robustness:* The second effect concerns the information capacity of the connections in the network. The high density of local connections in the regular network leads to synaptic redundancy, as activation between any two nearby neurons can take multiple paths between them. Necessarily, redundancy where more than one connection carries the same information results in a reduction in information capacity elsewhere in the network, as previously shown, but increased redundancy should also lead to networks which are more robust to damage.

To test this prediction, a profile of deletion (without com-

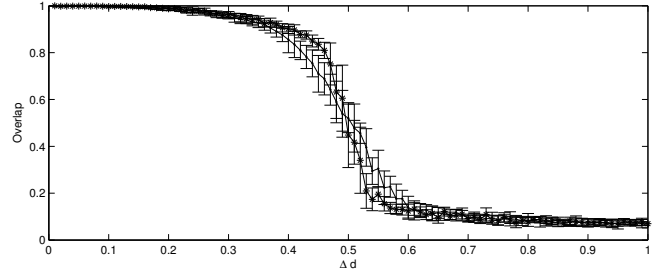


Fig. 8: Deletion without compensation in random (star markers) and highly-ordered (dot markers) networks.

pensation) was obtained for networks connected with small-world ($p(\text{rewire}) = 0.01$) and flat-random connectivities in 1600-unit networks with connectivity $K = 0.125N$. The resulting plot in figure 8 shows a marginally smoother rate of decline and greater longevity of performance in the small-world network (dot markers) than in the random network (star markers), indicating that the high local connectivity density does indeed lead to redundancy and hence greater robustness to damage, but at the expense of lower capacity. Nevertheless, the effects are relatively small overall.

D. Tau lesioning

To identify the changes in behaviour when more distributed, variable-rate tau damage occurs within the network, a network with $N = 800$ units was connected in a Gaussian manner and tau lesioning was performed according to the method described in section II-F, with random-set compensation. Two rates of tau lesioning were inspected: in addition to the standard rate in which the neuronal outputs were muted by an inverse Gaussian probability as a function of the neuron’s distance from the lesion centre, a second rate was tested in which the muting amount was squared so as to increase the speed with which the lesions resulted in full neuronal blocking, and the width of the distribution used for choosing new nearby tau lesion centres was doubled. Examples of the resulting comparable increase in lesioning can be seen in figure 9. A further, currently untested, method of altering the tau lesioning rate would be to consider each unit more than once until full blocking of all units occurs.

The results in figure 10 show a very different profile to basic deletion (see figure 2). Rather than a smooth decline in performance which tails off towards zero, a sudden catastrophic decline in performance occurs during a single step of tau lesioning. The performance then steadily recovers, but only to a fraction of the original performance, as the compensatory mechanisms attempt to ‘catch up’ with the sudden decrease in activation. As seen in figure 9, there are still areas of the network which are undamaged (transmission remains at 1), and it is likely that it is these areas which contribute to the above-zero final performance of the network.

Although the precise timing of the sudden decline varies randomly between test runs, each line on the graph traces essentially the same shape. Indeed, it was found that the differing rates of neuronal damping shown in figure 9 resulted

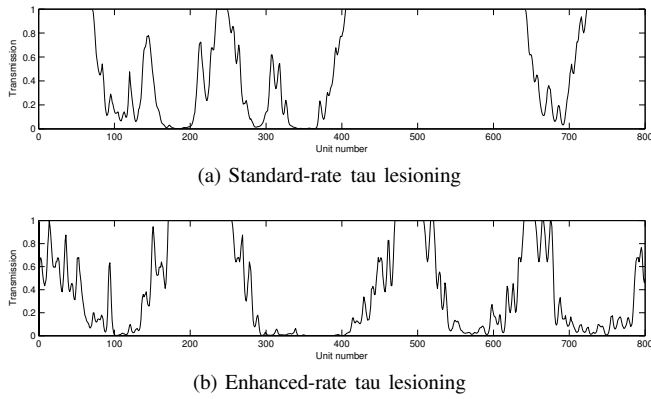


Fig. 9: Lesions applied to example networks after every unit has been considered once. In the second graph, the rate of reduction of neural transmission in each lesioning step has been squared and the horizontal spreading speed of the lesioning has been doubled.

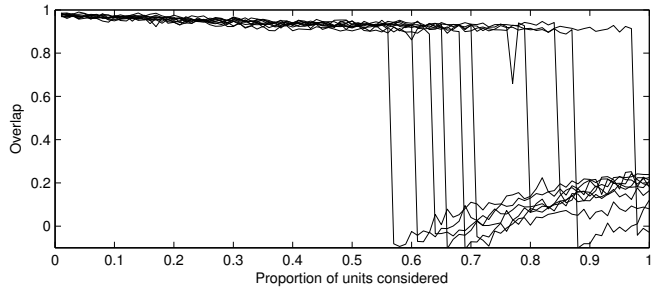


Fig. 10: Performance plotted against number of units considered for lesioning with tau (enhanced rate). Due to the large differences in the time of onset of impairment, individual test runs have been plotted for easier comparison.

in exactly the same profile of altered performance during lesioning, except that the performance drop-off was experienced correspondingly earlier or later with faster or slower tau lesioning rates.

To test that the observed gradient of learning in sets of patterns over time (figure 3) still occurs with tau lesioning, the experiment in section III-B was re-run in a Gaussian-connected network with tau lesioning instead of deletion. The patterns used for compensation were drawn at random from all those previously stored. The results in figure 11 are comparable with those in figure 3, indicating that tau lesioning in this way does not destroy the effect of reduced retrieval of recent compared to remote patterns.

IV. CONCLUSIONS AND FURTHER WORK

This work has presented updates to a long-standing and widely-cited computational model of Alzheimer’s Disease [23], including first successfully replicating, and then extending, the experiments of Horn et al. [13] on the effects of local, field-dependent synaptic compensatory mechanisms within the model.

The differing effects of using recent, remote, and random sets of memories to calculate compensatory signal terms

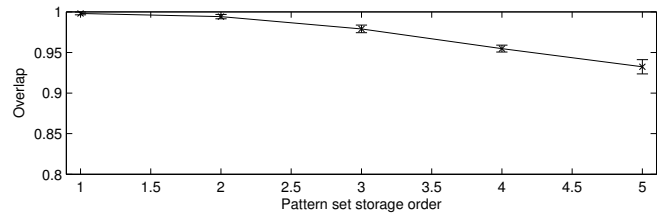


Fig. 11: Gradient of performance on sets of patterns stored at different stages of lesioning with tau.

has been shown, revealing that the network is sensitive to the choice of which set is used. Using only remote memories to calculate the signal term results in greater noise within the compensatory mechanism, and an earlier decline in performance as synapses are deleted (much closer to the 10 – 30% range seen in AD patients [18]). The implications for AD patients are shown in the context that initial retrieval of remote memories at early stages of damage is actually more reliable than with recent memories: if the brain makes use of this effect and uses the more readily-available remote memories to calculate compensation, not only do the recently-stored memories continue to become less reliable than the remote memories, but the noise in the system leads to earlier onset of catastrophic decline.

Speculatively, if the biological realisation of synaptic compensation via memory retrieval could be considered as dreaming (as postulated by Horn et al. [13]), these results are consistent with the idea that the greatest compensatory success is likely to be found by using memories and states acquired throughout an individual’s lifetime in the compensatory mechanism, or by using those memories most recently obtained, rather than primarily memories from early life. Further studies to examine any potential link between this effect and any reported fixation during dreaming on remote memories in AD patients (either prior to, or after, onset of symptoms) could yield important results.

It has also been shown that network capacity and resilience is related to the regularity of connections within the network. High small-world clustering coefficients lead to redundancy within the network, meaning greater resilience to damage but at the expense of lower capacity, as well as longer pattern retrieval times. This is consistent with the findings of Supekar et al. [27] who examined small-world functional networks in the brain and found a key correlation between loss of small-world connectivity and onset of AD symptoms. Further examination of the relationships between small-world clustering, robustness, retrieval speed and network capacity could be revealing, as well as studies into how this operates within the principle of neural Darwinism (pruning of weaker synapses during brain development).

Lesioning with simulated tau rather than standard synaptic deletion has been shown to create a very different profile of damage by allowing all neurons and synaptic connections to remain present (so output patterns are not artificially altered) and instead damping inter-neuronal transmission. Whilst initially offering a much more graceful decline in performance

due to the persistence of synaptic connections and output units, consistent with the slow degradation seen in AD, the drop-off in performance when it finally occurs is much more severe with tau lesioning than with synaptic deletion despite some later compensatory recovery of performance.

Further work will be needed to ascertain whether this deletion profile offers a more plausible explanation of AD symptoms and whether the observed temporary improvement in recall after some level of catastrophic damage can be medically corroborated, but it must be borne in mind that tau pathology represents only a subset of the processes underlying AD. It would be beneficial to extend this concept and show in a similar way the effects of alternative AD pathologies. Of particular interest are the beta-amyloid mechanism and its extension, the N-APP hypothesis, in which a fragment of the amyloid precursor protein (N-APP) is found to be capable of binding to the DR6 cell-death receptors of neuronal cell bodies and axons, with the effect of accelerating apoptotic cell death. The apoptotic mechanism involves the release of *caspases*, of which caspase 6 is capable of cleaving the N-APP fragment from existing β -amyloid deposits, leading to a cascade of neurodegeneration [19].

The Tsodyks and Feigl'Man model studied in this work is only a basic associative network with limitations in processing ability and a relatively constrained range of behaviour. More sophisticated artificial neural network-based models such as LEABRA [21], spiking neurons [9], and reservoir networks [17] are available and could provide further insights into the effects highlighted in this paper. In particular, it would be interesting to examine the effects of synaptic compensation and connectivity strategies within a reservoir computing framework due to the large potential for exploration of the currently poorly-understood dynamics, and the greater computational power (potentially offering the representation of more varied symptoms of AD than simple pattern recall) of these systems.

V. ACKNOWLEDGEMENT

With thanks to my academic supervisor, Dr. John Bullinaria, for his comments.

REFERENCES

- [1] L. Aakerlund and R. Hemmingsen. Neural networks as models of psychopathology. *Biological Psychiatry*, 43(7):471–482, 1998.
- [2] M. Abeles, E. Vaadia, and H. Bergman. Firing patterns of single units in the prefrontal cortex and neural network models. *Network: Computation in Neural Systems*, 1(1):13–25, 1990.
- [3] C. Ballatore, M.Y.L. Virginia, and J.Q. Trojanowski. Tau-mediated neurodegeneration in Alzheimer's disease and related disorders. *Nature Reviews Neuroscience*, 8(9):663–672, 2007.
- [4] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3):186–198, 2009.
- [5] W. Duch. Computational models of dementia and neurological problems. *Methods in Molecular Biology*, 401:305–336, November 2007.
- [6] M.R. Evans. Random dilution in a neural network for biased patterns. *Journal of Physics A: Mathematical and General*, 22:2103, 1989.
- [7] S.C. Feinstein and L. Wilson. Inability of tau to properly regulate neuronal microtubule dynamics: a loss-of-function mechanism by which tau might mediate neuronal cell death. *Biochimica et Biophysica Acta (BBA)-Molecular Basis of Disease*, 1739(2-3):268–279, 2005.
- [8] P.T. Francis, A.M. Palmer, M. Snape, and G.K. Wilcock. The cholinergic hypothesis of Alzheimer's disease: a review of progress. *Journal of Neurology, Neurosurgery & Psychiatry*, 66(2):137, 1999.
- [9] W. Gerstner and W. Kistler. *Spiking neuron models: An introduction*. Cambridge University Press New York, NY, USA, 2002.
- [10] J. Hardy and D.J. Selkoe. The amyloid hypothesis of Alzheimer's disease: progress and problems on the road to therapeutics. *Science*, 297(5580):353, 2002.
- [11] M.E. Hasselmo. Runaway synaptic modification in models of cortex: Implications for Alzheimer's disease. *Neural Networks*, 7:13–40, 1994.
- [12] R.E. Hoffman and T.H. McGlashan. Neural network models of schizophrenia. *Neuroscientist*, 7(5):441–454, 2001.
- [13] D. Horn, N. Levy, and E. Ruppín. Neuronal-based synaptic compensation: a computational study in Alzheimer's disease. *Neural Computation*, 8(6):1227–1243, 1996.
- [14] G.V.W. Johnson and W.H. Stoothoff. Tau phosphorylation in neuronal cell function and dysfunction. *Journal of cell science*, 117(24):5721, 2004.
- [15] M.D. Kopelman. Remote and autobiographical memory, temporal context memory and frontal atrophy in Korsakoff and Alzheimer patients. *Neuropsychologia*, 27(4):437–460, 1989.
- [16] W.B. Levy. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, 6(6):579–590, 1996.
- [17] M. Lukosevicius and H. Jaeger. Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3):127–149, 2009.
- [18] L. Minati, T. Edginton, M. Grazia Bruzzone, and G. Giaccone. Reviews: Current Concepts in Alzheimer's Disease: A Multidisciplinary Review. *American Journal of Alzheimer's Disease and Other Dementias*, 24(2):95, 2009.
- [19] D.W. Nicholson. Good and bad cell death. *Nature*, 457(7232):970–971, 2009.
- [20] A. Nikolaev, T. McLaughlin, D. O'Leary, and M. Tessier-Lavigne. N-APP binds DR6 to cause axon pruning and neuron death via distinct caspases. *Nature*, 457(7232):981, 2009.
- [21] R.C. O'Reilly. Generalization in interactive networks: The benefits of inhibitory competition and Hebbian learning. *Neural Computation*, 13(6):1199–1241, 2001.
- [22] B.H. Ridha, J. Barnes, J.W. Bartlett, A. Godbolt, T. Pepple, M.N. Rossor, and N.C. Fox. Tracking atrophy progression in familial Alzheimer's disease: a serial MRI study. *The Lancet Neurology*, 5(10):828–834, 2006.
- [23] E. Ruppín and J.A. Reggia. A neural model of memory impairment in diffuse cerebral atrophy. *The British Journal of Psychiatry*, 166(1):19–28, 1995.
- [24] E. Ruppín and J.A. Reggia. Patterns of functional damage in neural network models of associative memory. *Neural computation*, 7(5):1105–1127, 1995.
- [25] A. Savioz, G. Leuba, P.G. Vallet, and C. Walzer. Contribution of neural networks to Alzheimer disease's progression. *Brain research bulletin*, 2009.
- [26] T.L. Spire-Jones, W.H. Stoothoff, A. de Calignon, P.B. Jones, and B.T. Hyman. Tau pathophysiology in neurodegeneration: a tangled issue. *Trends in neurosciences*, 32(3):150–159, 2009.
- [27] K. Supekar, V. Menon, D. Rubin, M. Musen, and M.D. Greicius. Network analysis of intrinsic functional brain connectivity in Alzheimer's disease. *PLoS Comput Biol*, 4(6):e1000100, 2008.
- [28] P. Tiraboschi, L. A. Hansen, L. J. Thal, and J. Corey-Bloom. The importance of neuritic plaques and tangles to the development and evolution of AD. *Neurology*, 62(11):1984–9, June 2004.
- [29] M.V. Tsodyks and M.V. Feigl'Man. The enhanced storage capacity in neural networks with low activity level. *EPL (Europhysics Letters)*, 6:101–105, 1988.
- [30] D.J. Watts and S.H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.