

Computational modelling of Alzheimer's disease
RSMG 6 (progress report)

Mark Rowan
School of Computer Science
University of Birmingham
`m.s.rowan@cs.bham.ac.uk`

Supervisor:
Dr. John Bullinaria

Thesis group members:
Prof. Xin Yao (RSMG representative)
Dr. Jon Rowe

Additional advisor:
Prof. John Jefferys
(School of Neuroscience)

September 2012

1 Introduction

As a short reminder of the work I am undertaking, I am investigating mechanisms and progression of Alzheimer’s disease, including the relationship with cognitive decline, using artificial neural network models.

Early parts of my work investigated the effects of a mechanism known as *synaptic scaling* (or compensation), in which the inputs to neurons are amplified via a homeostatic mechanism to make up for reduced input due to death of neighbouring cells, but where this amplification and over-excitation can itself lead to accelerated neural death via calcium toxicity. The work makes use of information theory to characterise neural significance, according to my hypothesis that the “least significant” neurons are killed first as a result of the scaling process, thus explaining the delayed onset of cognitive symptoms.

Currently I am working in collaboration with Bill Lytton’s lab in Brooklyn, New York, to investigate these effects in greater detail in a biologically-realistic spiking neural network simulation with STDP learning.

2 Work done so far

2.1 Identification of candidate models

After the last thesis group meeting in January 2012, some time was spent exploring the Yale ModelDB¹ (Hines et al., 2004), which presents a comprehensive list of peer-reviewed computational models with accompanying source code covering various areas of neuroscience, including abstract connectionist models of neurological disorders such as Parkinson’s disease, highly detailed “multi-compartment” biophysical models of individual neurons, and networks of biologically-realistic neurons for potentially suitable biologically-realistic simulation models. The purpose of this search was to identify a model or models which could realistically be extended to simulate the synaptic scaling-driven pathology investigated as part of the Frontiers paper (Rowan, 2012).

Ten candidate models were identified from this search based on a number of criteria, and evaluated for their suitability to the task of extending each model to incorporate synaptic compensation and targeted neuronal deletion according to the compensation rates. A report of ≈ 4500 words in length was written to document this process, and some elements of this report will probably be incorporated into a relevant chapter of the thesis.

2.1.1 Neymotin et al. (2011c) model

As one of the two potentially most suitable models found during the database search, Neymotin et al. (2011c) present a large network simulating a section of neocortex containing 4230 point neurons across 14 classes, with excitatory AMPA and NMDA as well as inhibitory GABA-A receptors. The network is

¹<http://senselab.med.yale.edu/modeldb/>

arranged into 9 minicolumns giving 470 neurons per column. Connectivity data is taken from a variety of recent sources, although some estimates have had to be made, and data is taken from various different (albeit related) animal species. The model displays a full range of oscillation frequencies under background noise stimulation conditions, and benefits from an extensive analysis of the network’s firing dynamics via the technique of rewiring neurons to slowly alter the distribution and number of “hubs” within the network, as well as detailed information-theoretic analysis of the flow of data throughout the network (Neymotin et al., 2011a), for which source code is also available in the ModelDB. This particular study may provide useful mechanisms and benchmarks for analysing the effects of beta-amyloid lesioning on the information contribution of remaining neurons in the network.

The Neymotin et al. (2011c) model contains fixed synaptic weights between neurons, meaning that no associative learning of patterns is possible and initial information contributions for each neuron may be arbitrary, but another paper based on this model (Neymotin et al., 2011b) implements a STDP learning rule in the AMPA receptors of excitatory neurons to tune the oscillatory dynamics of the model. The authors predict that “homeostatic control mechanisms must balance learning at excitatory-to-excitatory and excitatory-to-inhibitory synapses” in order to avoid saturation of firing and transition to an epileptic state, matching the predictions of Turrigiano (2008) regarding the need for synaptic down-scaling during learning, and making a compelling connection to the work I have been doing on synaptic compensation (Rowan, 2012).

The model has been experimentally verified by the authors in data taken from the rat neocortex and shown to match physiological data with a high degree of accuracy, and the authors highlight the fact that their model shows oscillations similar to those of a ‘waking’ brain state, in specific contrast to the model of Traub et al. (2005) (one of the other candidate models examined) which shows a smaller range of oscillatory frequency peaks. The model benefits greatly from nearly real-time processing speed in the NEURON environment (i.e one second of simulated time takes approximately one second of processor time), which makes the model suitable for sufficiently long runs to simulate synaptic scaling (in the order of 24 hours of simulated time).

2.1.2 Cutsuridis et al. (2010) model

The other strong candidate model, this time of the CA1 region of the hippocampus, was presented by Cutsuridis et al. (2010). The model demonstrates associative storage and retrieval of patterns in the CA1 region of the hippocampus via a simulated STDP learning process. In a similar approach to Ruppin and Reggia (1995), the authors provide an “overlap” measure of recall performance as patterns are stored and retrieved using a noisy cue, thus allowing a quantifiable analysis of the network’s performance on a standardised task. The model investigates a theory by Hasselmo (Hasselmo et al., 2002) which predicts that theta rhythms (4-7Hz) dictate a mode switch between episodes of storage of new information and retrieval of old information.

STDP is not modelled biophysically at the molecular level, but a biologically-plausible localised mathematical learning rule defines the strengthening and weakening of relevant AMPA synapses based upon realistic simulations of the appropriate firing patterns within the network. In this particular model, only a subset of the synapses on the pyramidal (excitatory) neurons are modifiable during the STDP learning process, and none of the inhibitory synapses are modifiable.

In addition to the applied task of learning and retrieving patterns in an associative manner, the model provides realistic simulations of gamma and theta rhythms with verification according to data from the literature on anaesthetised rats, and the network wiring data is taken from a variety of experimentally-derived literature sources.

Neurons within the model are multi-compartment (13-17 separately modelled compartments), in comparison to the single-compartment point neurons of Neymotin et al. (2011c), and are presented in five distinct classes (one excitatory and four inhibitory). The model provides various Ca^{2+} and K^+ channels with AMPA, NMDA, GABA-A and GABA-B receptors, offering one of the most realistic simulations of all of the studied models. This means that there is a trade-off, which is that only 100 cells are modelled due to the additional complexity of the simulation. Even with this small network size, the model takes a long time to process a single second of simulation time. It is written to run in parallel using the NEURON environment, however, which mitigates some of the slow computation speed: using 20 CPU cores of the University of Birmingham's BlueBEAR computing cluster, 2s of simulated time takes approximately 200s to run, in comparison with approximately real-time running for the Neymotin et al. (2011c) model.

2.1.3 Choice of model

After evaluating these two models, as well as the other eight candidates from the ModelDB, it was decided to progress with the Neymotin et al. (2011c) model, primarily due to its biologically-validated connectivity which is vital for investigating the information-selectivity hypothesis, but also because of its vastly reduced computation time compared to the Cutsuridis et al. (2010) model (in which individual runs for the required simulation timespan would take weeks to complete, particularly if the network was enlarged to more than 100 cells).

2.2 Collaboration with SUNY Downstate, Brooklyn

After selecting the Neymotin et al. (2011c) model, I made contact with the author to enquire about the possibility of collaboration. The author, Sam Neymotin, is based in Bill Lytton's Neurosim lab at State University of New York, Brooklyn, and their team agreed to collaborate with me, giving me access to their source code and computing facilities, as well as answering questions as required to help me learn how the model operates.

Much time since then has been spent learning my way around the source code, and teaching myself the NEURON simulation environment as well as its associated programming language, *hoc*. The first contribution has been extension of the model to allow very long runtimes (over 27 hours of simulated time so far), and saving of this extended simulation data to disk, as well as Python software to create graphs and figures from the resulting files. The model has also been converted to work with the SoCS *cluster1* cluster.

Once this groundwork had been completed, the next task was implementation of synaptic scaling in ‘normal’ operating mode, so that the network can respond to reduced input due to death of neurons by scaling up the sensitivity of the AMPA receptors, as described by Turrigiano (2008) and Small (2008). Example plots of excitatory cell activity rates during deletion with and without synaptic scaling can be seen in figure 2.2.

2.2.1 Using scaling to balance STDP learning

As part of this collaboration with SUNY Downstate, the current intention is to write and submit a paper investigating the synaptic scaling effects when combined with STDP learning. This is to follow-on from work by Neymotin et al. (2011b) in which the network was trained using STDP from a training signal, with a resulting shift in the peak frequency of its natural oscillations. The network showed runaway potentiation of synapses during training with high-frequency (16Hz) signals due to a lack of compensatory down-scaling, but by using bidirectional synaptic scaling during learning, these effects may be mitigated (Turrigiano, 2008).

The paper is targeted for ICANNGA 2013, proceedings of which will be published in Springer Lecture Notes in Computer Science.

2.3 Teaching duties

During July and August, approximately four weeks were spent full-time on teaching preparation for the 3rd year Intelligent Robotics module. There will also be 4–5 hours of lab demonstrations per week in semester 1. I am currently awaiting confirmation of my teaching duties for the remainder of the year until end of April 2013. My teaching work for 2011-12 was recognised with a CLAD Teaching Award for “Best Teaching Assistant”, which was gratefully received on 27th June.

2.4 Courses attended and presentations given

Since January 2012 I have attended the Graduate School’s Academic writing course (15th February) and Viva preparation course (24th April). I gave a presentation to the Neuroscience department’s Neuronal Networks group on 29th February, from which I gained useful insights thanks to a conversation with Zsuzsa Nagy (see section 3.1). I also gave an updated version of this talk at the *3rd BEAR Postgraduate Conference* on 28th June.

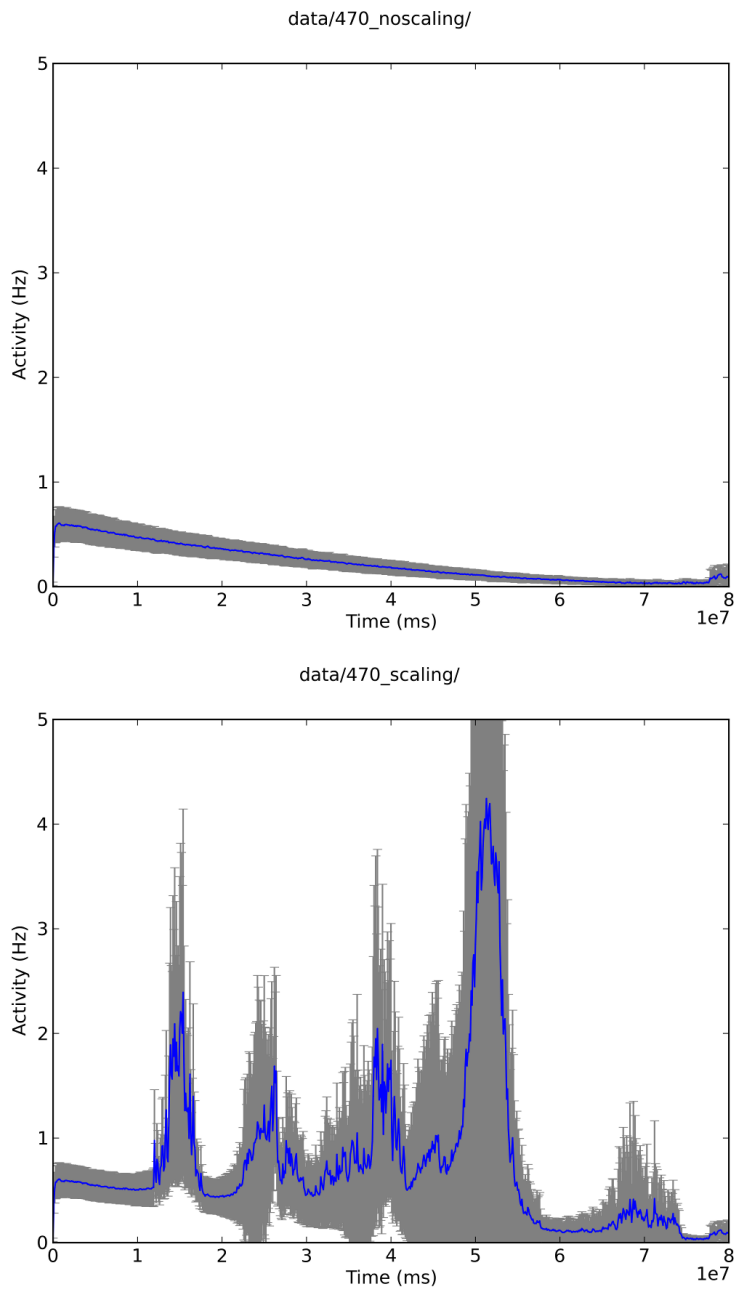


Figure 1: Average firing rates of excitatory cells as neurons in a network of 470 cells are deleted and external inputs are linearly reduced. In the second figure, synaptic scaling maintains activity for a longer period of time, but with several large peaks due to over-compensation.

Additionally I was invited by a potential postdoctoral employer to give a talk summarising all of my research at ETH Zürich on 30th August. Writing the talk was a very useful exercise in joining together the different strands of my research into a coherent story, in preparation for writing the thesis. I also gave a dry-run of the talk in the School of Computer Science before heading to ETH.

Finally, I attended the INCF Short Course on Neuroinformatics in Munich (7th–8th September), for which I was awarded a travel scholarship. Although it was only two days in length, the course provided valuable insights into the computational modelling process, and advice from other current researchers in the field including methods of biological validation and statistical analysis of neural spike trains.

3 New problems encountered

3.1 Beta-amyloid: cause or symptom?

Following on from the talk given to the Neuronal Networks group in the Neuroscience department, a conversation with Zsuzsa Nagy revealed that current medical opinion on beta-amyloid as a direct cause of Alzheimer’s disease is divided. Whilst Small (2008) states “the build-up of aggregated forms of Abeta leads to synaptic loss and to cognitive dysfunction”, other authors such as Joseph et al. (2001) describe alternative theories in which beta-amyloid is merely a prominent marker of some other underlying process: “Abeta has not proven to be both necessary and sufficient for the development, neurotoxicity, and cognitive deficits associated with this disease”. Even Hardy (2009), as the “father” of the amyloid hypothesis of Alzheimer’s disease, presents a critical reappraisal of the current state of the research, outlining a number of limitations in current knowledge and explanation of the disease mechanism.

This difference is more significant for clinical researchers than it is for my research, as I am primarily interested in the cascades of degeneration and progression and subsequent cognitive defects observed in Alzheimer’s disease, regardless of the (medical) cause, whereas clinical treatments focussed entirely on eliminating beta-amyloid may, at least according to Joseph et al. (2001) and Hardy (2009), be largely attacking the wrong area. However the suggestion that the progression theory (Small, 2008) could still explain changes in neural firing dynamics during Alzheimer’s disease and other “disconnection events” still appears to fit with current evidence.

As far as I can tell, this will not mean that my research direction will have to change in any way, but rather that I should simply be careful not to make direct claims that “beta-amyloid is the primary driver of Alzheimer’s disease”. By maintaining a separation of *causes* (unknown) and *effects* (what I am studying), it should be possible to avoid any controversy around this issue.

3.2 Information contribution measures

The Shannon mutual information measure used in the *Frontiers* paper (Rowan, 2012) is a useful basis for measuring the information contribution of specific neurons to the wider network, but in the framework of the Ruppin and Reggia associative neural network model (with storage of discrete patterns) the information measure essentially reduces to a simple count of the number of patterns in which the neuron fires, normalised by the total number of patterns. Clearly, this is not a highly-detailed measure of information contribution, and the paradigm of counting the number of discrete stored patterns does not easily apply to a biophysically realistic spiking network simulation.

A recently published method for measuring information contribution of individual neurons in a spiking network takes a highly-efficient approach whereby each neural spike is reduced to a delta function and then passed through a Fourier transform to find its covariance alongside every other neuron in the network (Crumiller et al., 2011). This method is able to determine the contribution of each neuron to every other neuron in the network for large groups of neurons.

Another method, used by Neymotin et al. (2011a), is Spectral Granger Causality (Kamiński et al., 2001), an extension of Granger Causality (Granger, 1969) designed to allow analysis of causality relationships across different frequency spectra (such as neural oscillations). It is not yet clear whether this method will be applicable to individual-neuron information calculations, or whether it can only be used for whole-layer calculations. A further measure used in the same work is Normalised Transfer Entropy, for which source code exists which interfaces with the model, and which may also be useful.

Quiroga and Panzeri (2009) provide a comprehensive review of various current methods of obtaining information contributions from neurons, but most of these methods require high-dimensional covariance calculations and are thus limited to groups of only a handful of neurons at a time.

4 Changes to the plan of research

Other than the addition of a short paper in collaboration with the SUNY Downstate Brooklyn lab on using synaptic scaling to balance STDP learning effects, there are no changes to the current research plan.

5 Draft summary of planned thesis

5.1 Introduction

- Setting the scene and statement of problem.
- Frame the research as an exploration of computation and information processing in unreliable biological hardware.
- State research questions and hypotheses.

- State contributions to knowledge.
- List of published papers.
- Introduce structure of thesis.

5.2 Background

- Medical background to Alzheimer’s disease (drawn from RSMG3 as well as the IJCNN and Frontiers papers).
- Provide motivation for solving the problem (i.e. early diagnosis is important, cognitive defects are not well understood, etc).

5.3 Related work

- Survey of other computational models, with critical assessment and relation to own work.
- Early connectionist models: Ruppin and Reggia, Hasselmo, Menschik and Finkel.
- Explain limitations of the early models (see Frontiers paper and RSMG5).
- Current biophysical modelling: summarise the current state of the field in comparison to early modelling (e.g. current simulation environments: NEURON, Genesis, PyNN, BRIAN, etc), rather than individual models.

5.4 Exploration of a recurrent associative network model

- Fully describe the model.
- Report on connectivity, compensation and tau experiments from IJCNN paper.

5.5 Information selectivity of disease mechanisms

- Explain meaning of ‘mutual information’ in neural context, and why it is used in neuroscience (e.g. we can use a standard ‘bits’ measure).
- Acknowledge that beta-amyloid may not be the direct cause of AD (see Zsuzsa Nagy’s comments, and Joseph et al. 2001) but that the disconnection hypothesis of Small, and my subsequent observations of the progress of cognitive decline, are still viable.
- Methods, results and discussion from Frontiers paper, starting with deletion without compensation (to show necessity of compensation in the brain).

- Discuss loading-level implications raised by Frontiers reviewer. Include thoughts on whether multiple redundant representations of stored states could lead to utilisation always being near capacity (i.e. as new patterns are stored, some of the the redundant copies of other states may be replaced, but total utilisation is always high) – see Barrett and van Rossum, 2008.
- The majority of relevant literature (Barrett and Van Rossum, 2008; Treves and Rolls, 1991; Leibold and Kempter, 2006) investigates theoretical capacity *limits* for various networks and architectures, but not *average* utilisation. ‘Capacity’ has also been investigated in the context of information processing bottlenecks (Marois and Ivanoff, 2005), but not in terms of number of discrete stored states.

5.6 Biophysically realistic spiking neural network simulation

- Explain benefits of biophysically realistic network simulations over associative Hopfield network models.
- Explain choice of Neymotin model from the ModelDB.
- Describe implementation of synaptic scaling (including findings such as that inhibitory neurons are scaled *down* as activity drops (Turrigiano et al., 1998; Rutherford et al., 1998; Chandler and Grossberg, 2012).
- Report on experiments using STDP with associated compensation to avoid self-reinforcing epileptic activity (ICANNGA 2013 paper).

5.7 Information selectivity in biophysical model

- Discuss how using highly stochastic spiking neural nets and moving away from the paradigm of “n stored patterns” (where, instead, every state is a “pattern” determined arbitrarily by the random weights) gives better resolution for the information measure.
- Discuss choice of information measure (Crumiller et al. (2011) or Spectral Granger Causality (Kamiński et al., 2001)).
- Also highlight how a biologically realistic network gives the opportunity to make biologically testable predictions regarding changes of behaviour .
- Results and discussion.

5.8 Conclusions and further work

- State hypothesis, and demonstrate precision, thoroughness, contribution, and comparison with other related work. What have I shown? etc...

- Summary conclusions.
- Restate contributions to knowledge.

6 Statement

I have read and bookmarked the School's current documents giving advice on writing up, presenting and submitting theses, as well as the University's on submitting theses.

Signed: 

7 Timetable for remainder of the research

<i>Date</i>	<i>Tasks</i>
September 2012	Finish tuning synaptic scaling mechanism for biologically realistic model. Enable STDP in the model and investigate suitability to solving the problems identified by Neymotin et al. (2011b). Write up findings for ICANNGA 2013.
<i>8th October 2012</i>	<i>Submission deadline for ICANNGA 2013.</i>
Oct – Nov 2012	Implement scaling-dependent lesioning and investigate Alzheimer's disease pathology, whilst looking for potential bio-markers.
Nov – Dec 2012	Write paper detailing findings, e.g. for Frontiers or PLoS One.
Christmas 2012	Begin writing-up according to plan in section 5.
April 2013	Submit thesis.
<i>30th April 2013</i>	<i>End of EPSRC scholarship funding period.</i>

Table 1: Proposed timetable of specific aims until thesis submission.

References

- A.B. Barrett and MCW Van Rossum. Optimal learning rules for discrete synapses. *PLoS Computational Biology*, 4(11):e1000230, 2008.
- B. Chandler and S. Grossberg. Joining distributed pattern processing and homeostatic plasticity in recurrent on-center off-surround shunting networks:

- Noise, saturation, short-term memory, synaptic scaling, and bdnf. *Neural Networks*, 2012.
- M. Crumiller, B. Knight, Y. Yu, and E. Kaplan. Estimating the amount of information conveyed by a population of neurons. *Frontiers in Neuroscience*, 5, 2011.
- V. Cutsuridis, S. Cobb, and B.P. Graham. Encoding and retrieval in a model of the hippocampal ca1 microcircuit. *Hippocampus*, 20(3):423–446, 2010.
- C.W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- J. Hardy. The amyloid hypothesis for Alzheimer’s disease: a critical reappraisal. *Journal of neurochemistry*, 110(4):1129–1134, 2009.
- M.E. Hasselmo, C. Bodelón, and B.P. Wyble. A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, 14(4):793–817, 2002.
- M.L. Hines, T. Morse, M. Migliore, N.T. Carnevale, and G.M. Shepherd. Modeldb: a database to support computational neuroscience. *Journal of Computational Neuroscience*, 17(1):7–11, 2004.
- J. Joseph, B. Shukitt-Hale, N.A. Denisova, A. Martin, G. Perry, and M.A. Smith. Copernicus revisited: amyloid beta in alzheimers disease. *Neurobiology of aging*, 22(1):131–146, 2001.
- M. Kamiński, M. Ding, W.A. Truccolo, and S.L. Bressler. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. *Biological cybernetics*, 85(2):145–157, 2001.
- C. Leibold and R. Kempter. Memory capacity for sequences in a recurrent network with biological constraints. *Neural Computation*, 18(4):904–941, 2006.
- R. Marois and J. Ivanoff. Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*, 9(6):296–305, 2005.
- S.A. Neymotin, K.M. Jacobs, A.A. Fenton, and W.W. Lytton. Synaptic information transfer in computer models of neocortical columns. *Journal of Computational Neuroscience*, 30(1):69–84, 2011a.
- S.A. Neymotin, C.C. Kerr, J.T. Francis, and W.W. Lytton. Training oscillatory dynamics with spike-timing-dependent plasticity in a computer model of neocortex. In *Signal Processing in Medicine and Biology Symposium (SPMB), 2011 IEEE*, pages 1–6. IEEE, 2011b.

- S.A. Neymotin, H. Lee, E. Park, A.A. Fenton, and W.W. Lytton. Emergence of physiological oscillation frequencies in a computer model of neocortex. *Frontiers in Computational Neuroscience*, 5, 2011c.
- R.Q. Quiroga and S. Panzeri. Extracting information from neuronal populations: information theory and decoding approaches. *Nature Reviews Neuroscience*, 10(3):173–185, 2009.
- M. Rowan. Information-selectivity of beta-amyloid pathology in an associative memory model. *Frontiers in Computational Neuroscience*, 6(2), January 2012. doi: 10.3389/fncom.2012.00002.
- E. Ruppin and J.A. Reggia. A neural model of memory impairment in diffuse cerebral atrophy. *The British Journal of Psychiatry*, 166(1):19–28, 1995.
- L.C. Rutherford, S.B. Nelson, and G.G. Turrigiano. Bdnf has opposite effects on the quantal amplitude of pyramidal neuron and interneuron excitatory synapses. *Neuron*, 21(3):521–530, 1998.
- David H. Small. Network dysfunction in alzheimer’s disease: does synaptic scaling drive disease progression? *Trends in Molecular Medicine*, 14(3):103 – 108, 2008. ISSN 1471-4914. doi: DOI:10.1016/j.molmed.2007.12.006.
- R.D. Traub, D. Contreras, M.O. Cunningham, H. Murray, F.E.N. LeBeau, A. Roopun, A. Bibbig, W.B. Wilent, M.J. Higley, and M.A. Whittington. Single-column thalamocortical network model exhibiting gamma oscillations, sleep spindles, and epileptogenic bursts. *Journal of Neurophysiology*, 93(4): 2194, 2005.
- A. Treves and E.T. Rolls. What determines the capacity of autoassociative memories in the brain? *Network: Computation in Neural Systems*, 2(4): 371–397, 1991.
- G.G. Turrigiano. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell*, 135(3):422–435, 2008.
- G.G. Turrigiano, K.R. Leslie, N.S. Desai, L.C. Rutherford, S.B. Nelson, et al. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, pages 892–895, 1998.