

Coincidence Detection: Towards an alternative to Synaptic Plasticity

Francis Jeanson (fjeanson@connect.carleton.ca)

Institute of Cognitive Science, Carleton University
1125 Colonel By Drive
Ottawa, Ont., Canada

Abstract

This paper introduces neurally plausible mechanisms for the basis of active memory and association making without synaptic plasticity. After discussing the predominant hebbian view for memory and learning in cognitive science, its limitations are presented from an energetic and temporal perspective. In particular, it is shown that synaptic plasticity cannot account for short-term mechanisms in memory and association making. The foundations of alternative mechanisms based on the spatiotemporal dynamic properties of neural groups are then developed and supported by existing work. Finally, how these alternate mechanisms can address the problem of retrieval is addressed along with future considerations.

Keywords: online cognition, synaptic plasticity, short-term memory, active memory, sensory memory, association, learning, spatiotemporal dynamics, neural groups.

Introduction

In 1949 Donald Hebb emitted a groundbreaking hypothesis: brains remember and learn by virtue of a simple mechanism in which connections between neural cells strengthen (Hebb, 1949). According to his proposal, this would occur when both the cell emitting a signal and the cell receiving the signal fire coincidentally. Since Hebb's original postulation a much greater understanding of the mechanisms leading to metabolic changes in synaptic plasticity has been achieved. In particular, we now know that "synaptic strengthening" is not manifested as a single metabolic process but occurs at a number of loci and timescales in both pre-synaptic and post-synaptic stages of cell communication (Debanne et al., 2003). The predominant form of synaptic plasticity is known as long-term plasticity and can last minutes to hours. Here, both 'strengthening' (long-term potentiation) and 'weakening' (long-term depression) occur via coincident firings which induce exocytosis and endocytosis of both NMDA and AMPA receptors on the post-synaptic membrane (review by Cooke & Bliss, 2006). On the other hand, short-term plasticity, which lasts only a few milliseconds to seconds, takes place at the pre-synaptic level in the form of a change in size, concentration, or efficiency in release of neurotransmitter vesicles in response to incoming frequency of stimulation (Zucker, 1989; Stevens & Wesseling, 1999).

The simple local characteristics of synaptic plasticity have increasingly attracted researchers in cognitive science for the development of sophisticated theories of perception, information coding, learning and memory. Indeed, since the accrued popularity of connectionist approaches in the 1980's, a principal objective in neural modelling for cognition, aside from architectural criteria and representational constraints, has been to focus almost exclusively on finding the set of adequate synaptic weights that will enable an artificial neural

network to perform a desired task (Rumelhart & McClelland, 1986). Because homogeneous network weights would most often result in unsatisfactory function, individual weight assignment must be accomplished for each connection between any two artificial neurones. Since the early 1980's, vigorous research efforts in finding effective ways for setting neural synaptic weights led to the discovery of the now popular back-propagation method, but also to evolutionary methods via genetic algorithms, and others (Rumelhart et al., 1986; Nolfi & Floreano, 2000). Although effective, many of these techniques either bear little semblance to biological processes or are not naturally applicable at time scales within the neural network's lifetime (such as in evolutionary methods).

Although fixed heterogeneous networks have proven effective in problems such as pattern completion, function fitting, multiple representation coding, and others, most cognitive tasks require complex pluripotent dynamics that allow agents to transcend simple reactive behaviours. Indeed, over the past decade it has become clear that a careful investigation of those neural properties that engender adaptation to high degrees of variability is necessary. In addition to its unsupervised qualities and biological validity, hebbian learning has been heralded as the key mechanism to learning. By adjusting the connection sensitivity between any two cells based on prior activity, it becomes possible for neural networks to store representations that did not exist before, build new associations, and generate new functions during their lifetime (Di Paolo, 2000; Cooke & Bliss, 2006).

Despite this success it appears that a "synapto-centric" view has developed in the neuroscientific and cognitive science communities with respect to neural mechanisms leading to memory and learning. Increasingly, there are reasons to believe that alternate processes could accomplish similar functions without incurring as high a metabolic, temporal, and perhaps even functional cost as that caused by synaptic plasticity. With respect to memory, O'Reilly and Munakata, for instance, suggested that an important distinction be made between 'weight-based memory' and 'activation-based memory' (O'Reilly & Munakata, 2000). While weight-based memory corresponds to classic entrenchment of stimuli within synaptic connections, activation-based memory relates to mechanisms by which stimulus patterns are sustained within recursive or bi-directional networks. Given the general consensus that forms of memory lie at the base of complex cognition, activation-based memory is thought to play a central role in higher-order processes, such as planning, problem solving and decision making, in that they can maintain online memories in a readily accessible and quickly updatable

manner.

Though I support this activation-based view, I intend to emphasize, to a larger degree than have O'Reilly and Munakata, the significant informational potential of spatiotemporal dynamics without a priori dependency on synaptic weighting. Here, I defend the view that non-metabolic mechanisms¹ based on firing dynamics in groups of neurones can not only account for forms of memory but also types of learning that are particularly relevant to sensory and short-term cognitive tasks and perhaps constitute the founding constructs of higher-order processes.

Cognitive Limits of Synaptic Plasticity

Two major arguments can be made against synaptic plasticity as the sole mechanism for memory and learning: (1) the level of energy required, and (2) the temporal constraint that the metabolic processes impose.

Energy Cost of Synaptic Plasticity

While the brain represents 2% of the body's mass it consumes approximately 20% of its energy. With such a high energy cost there is reason to believe from an evolutionary standpoint that energy economy would be highly favourable to the individual. Within the complex circuit of energy regulation in the brain, LTP and LTD have been shown to depend on the synthesis of NMDA and AMPA receptors which are regulated by brain-derived neurotrophic factor (BDNF) molecules (Minichiello et al. 2002). This regulation occurs via de production of CREB proteins by BDNF using adenosine triphosphate (ATP), the energy source derived from glucose crossing the blood/brain barrier. Although synaptic transmission alone also demands BDNF regulation, a diminished requirement for NMDA and AMPA regulation signifies a much reduced demand on ATP consumption by the nerve cell. Overall, a lower energy requirement suggests greater self-sufficiency in the case of low food availability when favouring active processes over expensive metabolic ones.

Temporal Limits

In addition to the energetic limitations that synaptic plasticity imposes, temporal constraints could constitute the most significant drawback. Synaptic change places a non negligible time constraint for cognitive performance. During the 1990's intracellular patch-clamp stimulation and recording techniques have made it possible to accurately measure temporal aspects of cell to cell interaction (Gerstner & Kistler, 2002). Bi and Poo demonstrated with great accuracy the variation of long-term synaptic potentiation (LTP) and depression (LTD) via stimulation of embryonic rat hippocampal neurones (Bi & Poo, 1998). After 60 pulses at 1Hz, an important increase in current flux was detected if the presynaptic cell fired just before the postsynaptic cell, while decrease occurred when firing in the postsynaptic cell preceded firing in the presynaptic neurone. Importantly, their results

further showed that sensitization of the synapses increases slowly over a period of minutes following firing. This considerable delay in synaptic efficacy suggests that short-term cognitive tasks cannot rely on LTP and LTD for rapid storage and rapid association making.

A faster form of synaptic plasticity is known as short-term plasticity. Via high frequency stimulation (HFS), Stevens and Wesseling provide a reasonable account for short-term pre-synaptic response latency of hippocampal neurones (Stevens & Wesseling, 1999). In particular, after stimulation at 9Hz for 15 seconds, synaptic response does not augment within the first half second (500ms) and reaches full potential only a second after high frequency stimulation is complete. Although this shows that the time taken for short-term plasticity is at a significantly lower order of magnitude than for long-term plasticity, these metabolic processes still does not appear to be fast enough to account for fast memory formation and association making. In sensory memory tasks, for instance, event related potentials (ERPs) can be used to detect mismatches between a new deviant stimulus from previous standard stimuli. In auditory tasks, this mismatch negativity (MMN) response can occur within 150ms of the presentation of the deviant stimulus (Sussman et al., 1998). The interference produced by the mismatch suggests that, at the very least, an echoic sensory memory trace can take place within 150ms. Arguably, this mismatch effect could be explained by simple disruption of the stored standard stimuli by the deviant. However, Titinen et al. have recently shown that MMN response varies with respect to the degree of contrast of the deviant, where greater contrast leads to shorter MMN response times (Titinen et al. 2002). This suggests that sensory features (here frequency) are encoded in less time than short-term plastic mechanisms seem to account for.

Beyond memory, current theories of perceptual categorization similarly favour the view that the formation of prototypes is mediated by a perceptual representation system (Schacter, 1990). This system more generally captures the underlying principle that MMN indirectly reports, that is, that prior exposure to a stimulus improves performance on subsequent presentations of a feature-sharing stimulus. Indeed, when exposed to unfamiliar visual patterns, healthy subjects show a propensity to correctly classify different patterns that share common features by virtue of the fact that previous stimuli allowed for the formation of abstract prototypes (Coutinho et al., 2010). The short time course during which feature integration for categorization takes place suggests that active, rather than synaptic, neural processes underly this mechanism.

In recent years, increasing evidence that active mechanisms play a critical role in cognitive capacities has emerged. Although these have traditionally been associated with working memory, Curtis and Lee have recently suggested that a kind of 'persistent activity' within cortical regions plays a role beyond mere working memory integration by also enabling planning and decision making (Curtis & Lee, 2010).

¹Or at least scarcely metabolic mechanisms.

Indeed, persistent activity is believed to subserve the integration and maintenance of spatial locations, object identities, words, sounds, and other sensory memoranda, as well as playing a key role in the reinforcement of value functions. Despite the accrued evidence for the existence of rapid integrating processes without synaptic plasticity, few proposals with neuro-biological plausibility have been made.

In the following, I introduce elementary mechanisms by which simple memories and associations can be biologically instantiated via rapid and low energy consuming processes. Furthermore, I support this hypothesis by drawing from important findings in the areas of temporal dynamics and neural network topology in the biological and theoretical neural sciences.

Spatiotemporal Processes

In recent years, growing attention has been dedicated to the spatiotemporal properties of neural networks. Perhaps one of the most comprehensive reviews of this work can be found in Buzsaki's "Rhythms of the Brain" (Buzsaki, 2006). Despite significant advancements in the observations and analysis of spatiotemporal phenomena in the brain, little is understood with respect to their informational potential and actual role in cognition. Amongst the better known, neural oscillations and synchrony have attracted greater attention within the neuroscience community since von der Marlsburg's postulation that feature binding can be solved via correlated firing between cell groups (von der Marlsburg, 1981). Soon after, experimental evidence has gradually uncovered a number of functional properties in synergetic group of neurones. Carr and Konishi, for instance showed that auditory stimuli in barn owls can initiate a neural response allowing the distinction of the precise location of a sound source within 3 to 5 milliseconds, suggesting that sensory encoding within synchronized populations of cells constitutes a mechanism that can perform faster than mere cell-to-cell encoding via firing rates (Carr & Konishi, 1990). Today, oscillations have become an important topic in a number of neurodegenerative diseases such as parkinson's and epilepsy (Brown, 2003). Thus, it appears that fine temporal tuning of firing cells (and sub-threshold cell activity) must be regulated within critical bounds so as to perform complex information coding operations without incurring cognitive and physiological disorders.

Within the past two decades, it has become increasingly clear that beyond mere oscillation and synchrony, the precise timing in neural populations of specialized cells may contribute much more to active processing than previously thought. In particular, the Dynamical Cell Assembly Hypothesis suggests that coincidence detection² of incoming spikes should enable a number of important information coding properties in virtue of the topological configurations of cells, their connection delays and other temporal properties (for review see Fujii et al., 1996). Recently, Izhikevich and Hop-

²Not to be confused with synaptic coincidence detection in hebbian learning.

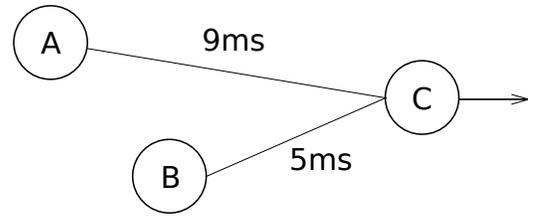


Figure 1: Coincidence Detection: A source neurone A connects to a target neurone C with a 9ms delay, and a source neurone B connects to target C with a 5ms delay. If neurone C requires the combined signals from A and B to fire, then C will only fire if B fires 4 milliseconds after A. C thus 'detects' the coincidence of incoming signals.

pensteadt suggested that spatially propagating wavefronts of propagating action potentials in sheets of coincidence detecting neurones, called neural groups, can in principle perform temporal signal processing, instantiate reverberating memories as well as solve logical functions (Izhikevich & Hoppensteadt, 2009).

A particularly important insight from this work is that the informational properties of neural groups would be particularly difficult to identify given the potentially immense variability in the firing pattern within a population of cells during a single stimulus. Moreover, the near chaotic dynamic properties of such networks suggests that no two presentations of the same stimulus will trigger the same neural response (Buzsaki, 2006). Hence, the development of biologically relevant neural network models should provide significant insight as to the informational capability of networks of coincidence detecting neurones.

In the following, I aim to introduce how such a network can, in principle, not only accomplish active memory storage but also establish active associations between initial stimuli without the use of synaptic plasticity.

Active Memory via Coincidence Detection

A neurone is a coincidence detector if it only fires when two or more incident signals trigger simultaneous post synaptic potentiations. The number of necessary incident spikes varies with respect to the activation threshold of the target cell. Hence, the time of firing of the source neurones is only relevant insofar as it affects the time of incidence of the incoming spikes on the target neurone. This timing is contingent on the firing time of the source neurones and the propagation delay of the signal to the target. For example, Figure 1 shows how a target cell C would only fire if action potentials (APs) from cells A and B arrived at the same time. For this to happen, cell A would have to fire 4ms ahead of cell B.

It is difficult to see the functional significance of coincidence detection when looking at a single cell. However, within the context of a large number of neurones the basis for a memory mechanism and association making takes shape. Figure 2 shows a group of coincidence detecting cells. Here,

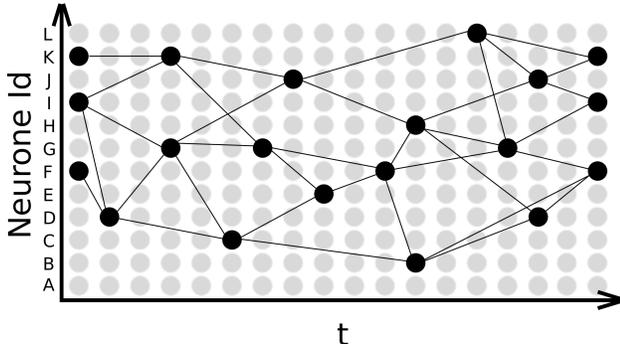


Figure 2: Neural population of 12 coincidence detecting neurons. With the initial activation of neurons F, I and K a firing chain is triggered within the population. All cells participating in the firing are part of the neural group. After 18 arbitrary time steps the original F, I, K neurons fire again. Thus, with the underlying connection topology and delays the simultaneous activation of F, I, K cells leads to a self-sustained firing loop.

we see a time plot of firing within a set of 12 cells (cell A to cell L). Although the topology is arbitrary and precise connection delays are unknown, we notice that if cells F, I, K fire simultaneously a chain of subsequent firing is initiated over time. Edges between cells show which connections are significant for a target cell to fire. The specific pattern of activation along with the cells that participate in this pattern constitute what is called a ‘neural group’ (Izhikevich, 2006). A particularly important aspect of coincidence detection neural networks is that slight timing differences at which inputs arrive in a population of cells will lead to the activation of a different neural group.

From Figure 2 one can see how given a different initial condition the activation of a much different neural group could form. In fact, each initial input may lead to a different neural group given that the probability of two connections having equal delay is sufficiently low. In a population of N cells the number of possible neural groups that are initiated by the simultaneous firing of at least one cell is $2^N - 1$. However, this assumes only the case where neural groups are formed via simultaneous input only. If the input can itself be asynchronous, i.e. input spikes arrive at different times, then the number of possible neural groups jumps drastically. Of course, the actual presence of neural groups will depend on the characteristics of the underlying connection delays and consequently on the topology of the network³.

Interestingly, if the connection topology of a given population of cells possess connection delays such that an initial input pattern is, after a period of time, re-triggered with the same temporal order, then the neural group is self-sustainable. Indeed, the neural group shown in Figure 2 illustrates such a

³Note that other factors, such as axonal size and degree of myelination, can also affect connection delays.

self-sustained neural group. Here, the initially active cells F, I, K trigger a firing sequence which overtime leads once more to the simultaneous firing of cells F, I, K. This suggests that the basis for simple, yet active, memories can take place within a population of coincidence detecting cells that possess appropriate connection characteristics. The statistical likelihood that self-sustaining groups are present in a given population of cells increases given that a significantly greater number of neural groups than individual cells exist in a given population, as shown above. This larger number of neural groups in proportion to cells is further supported by Izhikevich and his work on networks of spiking neurons with connection delays (Izhikevich, 2006). Thus, a subset of these neural groups are likely constituted by distinct self-sustaining groups; although the precise amount of self-sustained groups will depend on network and cell characteristics such as connection density, cell activation thresholds, connection delays etc.

Under appropriate neural and network conditions it now follows that a foundation for active memory processes can be instantiated within a network of coincidence detecting cells without the requirement for synaptic plasticity⁴. Still, questions such as *how much information can be stored given a neural group?* and *how do neural groups get selected for particular storage?* remain unanswered. These will need to be addressed with further investigation. Beyond memory, however, the ability for coincidence detecting cell populations to account for active forms of association making is also possible following similar principles. I introduce a potential mechanism in what follows.

Active Associations via Coincidence Detection

Upon the presentation of two stimuli within a narrow time period, two separate neural groups will form. If these two groups lead to a spatiotemporal pattern that triggers the formation of an available *self-sustained group* then an active association of both original stimuli is formed. That is, if both groups interact at a spatiotemporal location where a self-sustained pulsating chain is triggered then this newly formed self-sustained group constitutes an association of the two originating groups, and thus of the two originating stimuli. This association will subsist so long as it meets no significant perturbation. Figure 3 illustrates how two independent neural groups (group {ABCD} and group {IJKL}) that correspond to individual input stimuli can be associated if their spatiotemporal patterns happen to trigger a self-sustained neural group, e.g. group {EFGH}.

The fact that associations are only made contingent on the spatiotemporal availability of self-sustained groups reflects the *opportunistic* nature of this simple association making mechanism. Fitting well with this, there are reasons to believe that neural groups will likely interact given what we know

⁴Izhikevich does show, however, that the application of plastic rules within a population of coincidence-detecting neurons allows for a great degree of exploration into the self-organized properties of neural groups (Izhikevich, 2006).

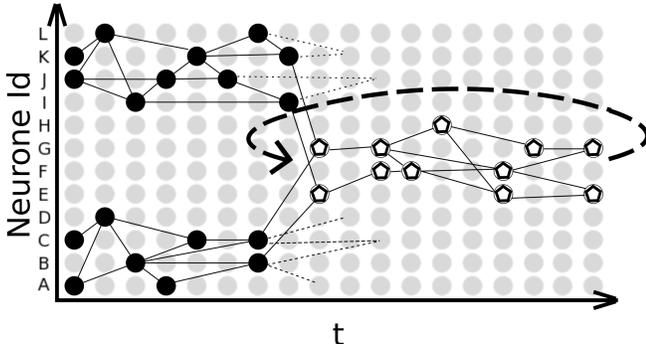


Figure 3: Associative neural group triggered by two independent input neural groups. Neural groups $\{ABCD\}$ and $\{IJKL\}$ (dark cells) possess cells that happen to trigger, in virtue of their spike times, the firing of self-sustained group $\{EFGH\}$ (segmented cells). Firing group $\{EFGH\}$ thus constitutes the active association of both stimulus induced groups $\{ABCD\}$ and $\{IJKL\}$. The dashed looping arrow shows that group $\{EFGH\}$ resets after a number of time steps resulting in a self-sustained group.

about the properties of brain connectivity. Indeed, the apparent small world network properties of the central nervous system suggest that there is a low degree of separation between any two neurones (Watts, 1999; Buzsaki, 2006). Hence the probability that two neural groups can jointly trigger a self-sustained group is relatively high. Although the possibility for such groups to interact due to their spatiotemporal properties is probable, this process requires that self-sustained neural groups exist in sufficient quantities for associations to be maintained. Fortunately, the small world hypothesis states that the brain is primarily composed of highly interconnected clusters with few inter-cluster connections. If cluster connectivity is sufficiently high then a large number of self-sustained neural groups can potentially form. Further investigation via neural simulation could help narrow down the actual topological and connective properties that a neural network must possess to sustain these processes.

The mechanism proposed here suggests that associations can be made without changing synaptic weights. Hence it seems that some preliminary form of learning can arise without short-term or long-term plasticity. Despite this, a number of principles related to learning are not obviously attainable by using this approach. In the following I discuss how coincidence based active memories and associations can be used within the larger context of the nervous system.

Contextual Application

A central concern with the solutions to active memory and association making provided above is *retrieval*. How can memories be retrieved for further cognitive use and how can an association be retrieved given the presentation of one of the two original stimuli? To address these, it is important to place active processing via neural groups into perspective.

One of the major advantages of this approach is that whenever an arbitrary firing pattern triggers a cascade of action potentials in a population of cells a memory can be retrieved as soon as a subset of this originating pattern fits a spatiotemporal pattern that is common to the memory. For example, the memory instantiated by the neural group in Figure 2 can be re-activated if, for instance, some ambient firing triggers cell C at time step t , cell G at $t+1$, and cell J at $t+2$ (start with $t=6$ in Figure 2). Of course, re-activation is an incomplete form of retrieval, however, by seeing how self-sustained neural groups can be embedded within larger neural groups, it should be possible for such memories to play the roles of pattern generators for further downstream processing. In addition, these active loops could serve as efficient mechanisms for the entrenchment of longer term memories by facilitating synaptic plasticity. The proposal that active memories can be instantiated within a population of coincidence detecting cells thus makes the prediction that specialized neighbouring structures and cells must be present so to integrate projecting outputs from the clusters of coincidence detection cells.

With respect to the retrieval of associations, a similar mechanism of re-activation can be applied. Retrieval typically implies that if two signals $S1$ and $S2$ are encoded in signal $S3$, then there must be a mechanism by which $S1$ can be retrieved from $S3$ given $S2$, and $S2$ can be retrieved from $S3$ given $S1$. In the approach proposed here, retrieval is possible if there exists joint projections from groups $\{ABCD\}$ and $\{EFGH\}$ that match a spatiotemporal configuration to re-activate the original stimulus group $\{IJKL\}$. This would require that some organizational processes guarantee that newly formed associative groups also possess re-entering projections to the source stimulus neural group in such a way that they are sufficiently tuned to participate in this retrieval process. This approach thus further predicts that a high degree of recurrent connections between clusters of coincidence detection cells must be present for active association making to occur.

Conclusion

Overall, the mechanisms of active processing presented herein offer a rather novel way of solving the problem of basic information storage and coding. By enabling the nervous system to ‘opportunistically’ establish new constructs from sensory signals and motor feedback it removes the conceptual hurdle of having to look for *intentional* or *attentional* processes that seek to organize information according to some overarching rules. Instead, the rules of organization from this perspective are decentralized local phenomena that take advantage of the temporal and topological configuration of their environment. Hence, by further exploring the biological conditions under which these mechanisms may take place, it should be possible, via simulation and analysis, to determine the extent to which coincidence detecting neurones can account for cognitive faculties.

Acknowledgments

I would like to thank Robert West and Tony White for their support and comments in preparing this paper and for guiding me in the experimental tests that led to the insights of this work (to be published).

References

- Bi, G. & Poo, M. (1998). *Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type*. *Journal of Neuroscience*, 18(24):10464-10472.
- Brown, P. (2003). *Oscillatory Nature of Human Basal Ganglia Activity: Relationship to the Pathophysiology of Parkinsons Disease*. *Movement Disorders*, 18(4): 357-363.
- Buzsaki, G. (2006). *Rhythms of the Brain*. Oxford University Press.
- Carr, C.E. & Konishi, M. (1990). *A Circuit for Detection of Interaural Time Differences in the Brain Stem of the Barn Owl*. *Journal of Neuroscience*, 70(10): 3227-3246
- Cooke, S.F., Bliss, T.V. (2006). *Plasticity in the human central nervous system*. *Brain* 129 (Pt 7): 1659-73.
- Coutinho, M.V.C., Couchman, J.J., Redford, J.S., Smith, D.J. (2010). *Refining the visual-cortical hypothesis in category learning*. *Brain and Cognition*, 74: 88-96.
- Curtis, C.E., Lee, D. (2010). *Beyond working memory: the role of persistent activity in decision making*. *Trends in Cognitive Sciences*, 14(5): 216-222.
- Debanne, D., Daoudal, G., Sourdet, V., Russier, M. (2003). *Brain plasticity and ion channels*. *Journal of Physiology, Paris*. 97(4-6):403-14.
- Di Paolo, E. A. (2000). *Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions*. *Proc. of the Sixth Inter. Conf. on the Sim. of Adap. Behav*. Cambridge MA: MIT Press.
- Fujii, H., Ito, H., Aihara, K., Ichinose, N., Tsukada, M. (1996). *Dynamical Cell Assembly Hypothesis - Theoretical Possibility and Spatio-temporal Coding in the Cortex*. *Neural Networks*, 9, 8:1303-1350.
- Gerstner & Kistler (2002). *Spiking Neuron Models. Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002.
- Hebb, D.O. (1949). *The Organization of Behavior*. Wiley, New York.
- Izhikevich, E.M. (2006). *Polychronization: Computation with Spikes*. *Neural Computation* 18, 245-282.
- Izhikevich, E.M. & Hoppensteadt, F.C. (2008). *Polychronous Wavefront Computations*. *International Journal of Bifurcation and Chaos*, 19(5): 1733-1739.
- Minichiello, L., Calella, A.M., Medina, D.L., Bonhoeffer, T., Klein, R., Korte, M. (2002). *Mechanism of TrkB-mediated hippocampal long-term potentiation*. *Neuron*, 36(1):121-37.
- Nolfi, S., Floreano, D. (2000). *Evolutionary Robotics: the biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- O'Reilly, R.C. & Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.
- Rumelhart, D.E., Hinton, G.E., Williams, R.J. (1986). *Learning representations by back-propagating errors*. *Nature*, 323:533-536.
- Rumelhart, D.E., J.L. McClelland and the PDP Research Group (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press.
- Schacter, D.L. (1990). *Perceptual representation systems and implicit memory: Toward a resolution of the multiple memory systems debate*. *Annals of the New York Academy of Sciences*, 608: 543-571.
- Stevens, F.C., Wesseling, F.J. (1999). *Augmentation Is a Potentiation of the Exocytotic Process*. *Neuron*, 22: 139-146.
- Sussman, E., Gomes, H., Noursak, J.M.K., Ritter, W., Vaughan, H.G. (1998). *Feature conjunctions and auditory sensory memory*. *Brain Research*, 793: 95-102.
- Tiitinen, H., May, P., Reinikainen, Naatanen, R. (2002). *Attentive novelty detection in humans is governed by pre-attentive sensory memory*. *Nature*, 372: 90-92.
- von der Marlsburg, C. (1981). *The correlation theory of brain function*. (Internal Report 81-2). Goettingen: Max-Planck Institute for Biophysical Chemistry.
- Watts, D.J. (1999). *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press.
- Zucker, R.S. (1989). *Short-term synaptic plasticity*. *Annual Review in Neuroscience*, 12: 13-31.