

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Period Doubling as a Means of Representing Multiply Instantiated Entities

Permalink

<https://escholarship.org/uc/item/9991f768>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 20(0)

Author

Sougne, Jacques

Publication Date

1998

Peer reviewed

Period Doubling as a Means of Representing Multiply Instantiated Entities

Jacques Sougné (J.Sougne@ulg.ac.be)

University of Liège
Department of Psychology
Bât B32, Sart Tilman
4000 Liège Belgium

Abstract

The problem of multiple instantiation is the ability to handle different instances of a unique object at the same time. For connectionist models that do not use a working area containing copies of items from a long-term knowledge base, the problem of multiple instantiation is a particularly difficult one. While people are able to deal with multiple instances, their performance when doing so is nonetheless poorer, which is not the case for symbolic models. A cognitive model should reflect competence, as well as its limits. Some connectionist solutions to the problem of multiple instantiation are mentioned in this paper. A new solution which makes use of semi-distributed representations is presented. This model does not separate the long term knowledge base from a working area and has no recourse to copies. This solution limits the process of multiple instantiation in a way that should better reflect human data.

Introduction

Multiple instantiation involves the simultaneous use of the same parts of the knowledge base in different ways. If you hear that "John is in love with Louise" and that "Louise is in love with John", you can easily infer that they should be happy. To do this inference, you must instantiate the predicate "is in love with" and the objects "John" and "Louise" twice. Precisely how this is done is the problem of multiple instantiation or "the type-token problem" (Norman, 1986; Dyer, 1991).

Symbolic models that load copies of pieces of knowledge into a working area before transforming them do not have any problem with multiple instantiation. They simply place several copies of the same content from the long-term knowledge base (LTKB) into the working area. However, for connectionist models, which do not use this copying process, multiple instantiation poses a serious problem. How can the same part of the knowledge base be related to different roles at the same time without making several copies of the knowledge in question? Multiple instantiation is even a greater problem for distributed representations. Two closely related concepts will, in principle, share nodes. If both concepts are needed simultaneously, their shared nodes must be instantiated twice.

An ability to handle multiply instantiated concepts assigned to different roles at the same time is required for many cognitive tasks. Some examples include:

- Transitive inferences: Knowing that Mary is older than Francis and that Francis is older than Jack, a cognitive system should be able to infer that Mary is older than Jack. This task requires two instantiations of the same predicate and two instantiations of Francis, each assigned to two different roles, "older object" and "younger object".
- Symmetric and non-symmetric inferences: From "John loves Louise" and "Louise loves Gray", the system should infer that "John is jealous of Gray." Here again, the task involves two instantiations of the predicate and two instantiations of "Louise", once in the role of "lovee" and once in the role of "lover".
- Recursion: Understanding this sentence: "The boy who hit the girl who hit the cat was my friend" requires two instances of the predicate "hit" and the concept "girl".

Connectionism and Multiple Instantiation

Classically, in a connectionist network there is no separation between LTKB and a temporary store (or a working area), in which copies of pieces of LTKB are loaded before transformation. In these models, activation of the LTKB creates a Short Term Memory (STM). For systems that do separate LTKB and STM (most traditional AI models), multiple instantiation is not a problem since the system can make as many copies of LTKB information as needed in STM. Without this copying process, neural nets suffer from "crosstalk." (Feldman, 1982). Adding "John loves Mary" to "Gary loves Rita" can lead to pseudo-memories (Dyer, 1991) like "John loves Rita". Even if we assume that John and Gary are correctly bound to the role of lover, and Mary and Rita to the role of lovee, both men and both women remain bound to the same respective roles. The system needs to distinguish the two facts by separating the two identical predicates and their respective roles bindings.

The problem of multiple instantiation arises in localist networks if two instantiations differ by more than one arguments value. For example, "Jack eats eggs and Jack eats fish" does not require separate instances of the predicate "eats" since this statement can be reduced to "Jack eats eggs and fish". However, when two sets of two items must be bound to identical pairs of roles, the system must be able to handle two copies of the predicate and argument slots. For example, "Jack eats eggs and Mary eats fish" cannot be reduced to "Jack and Mary eat eggs and fish," otherwise one cannot distinguish who eats what.

The problem for distributed representations is even more difficult. Multiple instantiation problems appear as soon as one node must be shared by entities that have to be differentiated. If an n-ary predicate must be represented where either predicate roles or their fillers need to share a common node, this node will have to be linked to different entities. In systems with distributed representations, the loss of a single node is of minor importance and, consequently, the problem will be a function of the proportion of shared nodes.

Relevance for Cognitive Science

Norman (1986) wondered if it was really necessary to solve the problem of multiple instantiation. The connectionist limitations involving multiple instantiation could be considered a virtue since humans have difficulties with tasks involving multiple instantiation. Empirical evidence can be found in psychological studies of reasoning (Sougné & French, 1997; Carreiras & Santamaria, 1997), of similarity and working memory (Baddeley, 1966), and of repetition blindness (Kanwisher, 1987; Morris & Harris, 1997). These studies show that multiple instantiation can indeed cause problems for humans. But they also show that the cognitive apparatus possesses the means to deal with it. Confronted with multiple instantiation people tend to be slower or to make more mistakes. A cognitive model should not only be able to deal with multiply instantiated concepts, but should also reflect human performance including difficulties (see Sougné, 1998).

Connectionist Solutions

There are three main types of connectionist solutions. The first uses two systems, one for the LTKB and another as working area where elements of LTKB are loaded. The second makes several copies of the same elements in LTKB. The third is the present attempt to solve the problem with different frequencies of oscillation.

Multiple Copies Loaded in a Working Area

Bookman and Alterman (1991) combine a localist semantic network that stores dependencies between concepts with a distributed network of semantic features that determines which schema slots will get filled. Each combination of a concept instance and its associated role will lead to a new schema. Another model, ABR-Composit of Barnden (1994) uses two systems, a Long term memory (LTM) and a Working memory (WM), both systems being connectionist networks. In this model, WM is composed of several registers which are filled with activation patterns from LTM.

Models that separate "LTM store" and "WM store" inadequately reflect difficulties people have when they perform multiple instantiation. For these models, even if WM store has a limited capacity, it is as easy to load one copy as to fill WM with copies of the same content from LTM unless this is prevented in an "ad hoc" manner. Other solutions have been developed, however, in which WM store is the activated part of LTM. These models are much

better at reflecting not only the ability, but also the difficulty, that humans have in doing multiple instantiation.

Multiple Copies of Concepts inside the LTKB

ROBIN (Lange & Dyer, 1989) separates roles from concepts, each concept has an associated node that outputs a particular constant value (called its signature). When a role node has the same activation as that of a concept signature, this concept is bound to the role. Multiple instantiation is performed by adding activations and signatures (see Lange, 1992).

SHRUTI (Mani & Shastri, 1993) uses synchrony of node-firing to bind objects to their roles. Multiple instantiation is achieved by the use of a bounded (usually 3) set of copies or banks of predicates and their argument slots and activation is directed to an uninstantiated copy by means of a switch. This model makes psychological predictions about both WM span and multiple instantiation abilities. SHRUTI predicts that the number of instantiations is limited and that the time required for doing multiple instantiation is proportional to the number of predicate banks.

Period Doubling

INFERNET (Sougné, 1996; Sougné & French, 1997) achieves variable binding through temporal synchrony of node firing. In short, when one node fires in synchrony with another, they are temporarily bound together. It has a limited WM span and the content of WM is maintained by oscillations. Once a node is activated, it tends to fire rhythmically at a particular frequency. It achieves multiple instantiation by means of period doubling. Nodes pertaining to a doubly instantiated concept will sustain two oscillations. This means that these nodes will be able to synchronize with two different sets of nodes. The following section describes the proposed solution in more detail.

INFERNET

INFERNET is a connectionist model using integrate-and-fire nodes. Each concept is represented by a cluster of nodes firing in synchrony. Concepts are bound to their roles by synchronous firing. Similar use of synchrony can be found in Shastri & Ajjanagadde (1993); Hummel & Holyoak, (1997); Henderson, (1996). For example, to represent the fact "John loves Louise", nodes belonging to "John" must fire synchronously with nodes belonging to "Lover" (Figure 1).

There is considerable neurobiological evidence for considering synchrony as a possible binding mechanism in the brain (see Roelfsema, Engel, König, & Singer, 1996; Singer, 1993).

Since concepts are represented by a set of nodes, INFERNET focuses on the distribution of node-firing times. If the firing distribution is tightly concentrated around the mean, the concept is considered to be activated.

There is neurobiological evidence (see Engel, Kreiter, König, & Singer, 1991) that if several objects are present in a scene, several groups of cells fire in distinct windows of

synchrony. In INFERNET, discrimination is achieved by successive windows of synchrony. Predicates and roles are linked by a specific temporal order. The activation of a predicate is always followed by the successive activation of its different roles, each of which is assigned to a particular window of synchrony.




Figure 1: Synchronization as a binding mechanism. Understanding “John loves Louise” requires grouping “Lover-John” and “Lovee-Louise” and discriminating them.

A number of neurobiological parameters are involved in a representation that relies on clusters of nodes firing simultaneously. The first is the frequency of oscillation. Some specific oscillatory activities seem to facilitate synchronization (Roelfsema et al., 1996; Singer, 1993). In INFERNET once a node is activated, it tends (but not necessarily) to begin oscillating at a γ frequency range, whose lower limit is 30Hz and upper limit varies according to various authors from 70Hz (Abeles, Prut, Bergman, Vaadia, & Aertsen, 1993) to 100 Hz (Wilson & Shepherd, 1995). The temporal gap between 2 spikes of a node is therefore from 10-14 to 33 ms. These γ waves have been observed to be associated with attention (Wang & Rinzel, 1995) and with associative memory (Wilson & Shepherd, 1995) and seem to be the best candidate for enabling synchronization and binding (Singer, 1993). The second key parameter is the precision of the synchrony at this frequency range. According to Abeles and al. (1993), this precision is about 5 ms, sometimes less, and depends on the frequency of oscillation. This allows us to approximate the number of windows of synchrony that can be differentiated, i.e., $25/5 = 5$, based on a typical frequency of 40Hz. If we assume that a window of synchrony corresponds to an item, a word, an idea, an object in a scene, or a chunk in working memory (WM), this puts WM span at approximately 5, with a small amount of variance since precision is proportional to oscillation frequency. This corresponds to estimates of human WM span (Cowan, 1998). The more the system needs to discriminate objects in WM, the more precise the synchrony should be. Since this parameter is bounded, it can lead to WM overload in which windows of synchrony are no longer distinguishable. Therefore, the number of distinct items and the number of predicate arguments in WM is limited (Sougné, 1996). Finally, the representation is maintained in WM by bursts of γ waves. Similar explanations for the brain’s ability to store short-

term memory items can be found in the literature (Shastri & Ajjanagadde, 1993; Lisman and Idiart 1995).

Inference with Multiple Instantiation

INFERNET uses a two-step process for drawing inferences. The first is to encode premises by temporarily “learning” the binding of objects to their respective roles. For example, Figure 2 shows the activity of the premises “John loves Louise” and “Louise loves John.” The second step is the network’s response to a query. For example, Figure 3 shows the activity of the concepts comprising the query, “Whose love is reciprocated?”

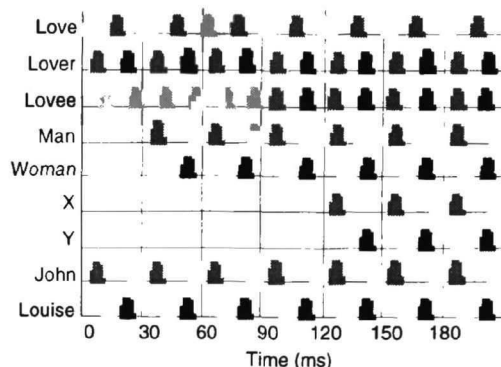


Figure 2: Concepts firing following the presentation of premises: “John loves Louise” and “Louise loves John” after a certain amount of learning. Vertically aligned histograms (denoting binding) have in the same color.

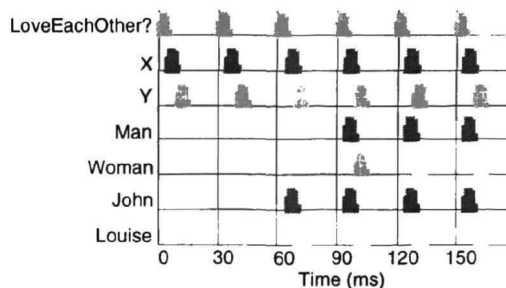


Figure 3: Concepts firing following the query: “Who loves each other?” once premises in B have been encoded.

During the encoding phase, the system is sensitive to synchronous firing of nodes. When two nodes fire in synchrony, if there is a connection between them, the strength of this connection will be positively increased and the delay associated will be adjusted. After learning, the firing of one of these two nodes will actively participate to the synchronous firing of the other. In short, this “learning phase” independently reproduces the synchronous firing of nodes detected from the input. The modifications of connection parameters decay over time to ensure that the system will be ready to new information.

INFERNET has a Long Term Knowledge Base that is used for encoding premises and answering queries. Figure 4 shows the knowledge necessary to make inferences about

love and jealousy. Arrows represent connections; they are tagged with numbers that indicate the time required to propagate activation. Specifically, in this example, a delay of 30ms corresponds to the lag between two spikes of a node oscillating at 33Hz. This delay ensures that these concept-node spikes will synchronize after 30ms. INFERNET also implements AND-gates, which require all inputs to reach the target at the same time. This is achieved by a set of excitatory and inhibitory links combined with presynaptic inhibition and facilitation (see Hawkins, Kandel, and Siegelbaum, 1993 for neurobiological counterpart). Unlike most links, these latter links act on *connections* rather than nodes (French, 1995; Shastri & Ajjanagadde, 1993). Similarly, XOR-gates are only on when one of the inputs is active and NOR-gates are only active when all inputs are silent. These gates are related to the neurobiological phenomenon of *coincidence detection* (see Konnerth, Tsien, Mikoshiba, & Altman, 1996).

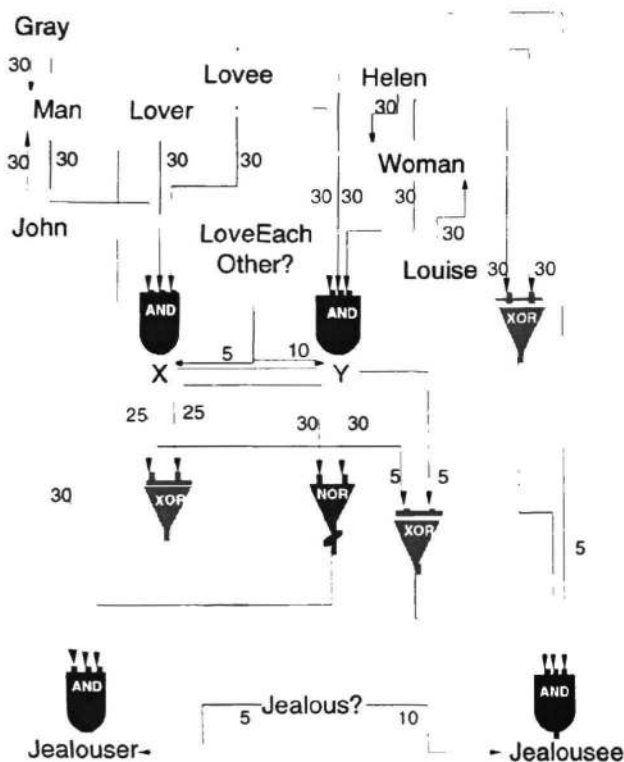


Figure 4: Connections, delays and gates required for reasoning about love and jealousy

The knowledge encoded, as shown in Figure 4, can correctly answer the query “Who loves each other?” and “Who is jealous of whom?” for all possible combination of two premises. In short, the connections of Figure 4 represent the following facts: People’s love is reciprocated if one individual is a woman, who is both lovee and lover, and the other is a man who is also both lovee and lover; a “jealouser” is a lover whose love is not reciprocated and the person whom she/he loves, loves someone else; a

“jealousee” is the person who is the lovee of someone who is loved by someone else.

During the premise-encoding phase, connection weights and delays will be modified by a Hebbian learning rule to reproduce synchronies. In Figure 2, “John” is synchronized with both “Lovee” and “Lover”. The connection strengths between “John” and these two roles will be increased. At a particular moment (95ms in Figure 2), the connections from “John” will be sufficient for the role of “Lovee” to be activated whenever “John” is activated. Thereafter, the expected activation of “Lovee” at 100ms will be prevented due to the refractory period of the “Lovee” nodes.

When the query comes (Figure 3), “Love-each-Other?” will be followed by the firing of “X” then “Y.” Since the strength of the connections between “John” and “X”, and between “Louise” and “Y” has increased, “John” and “Louise” will also fire.

When doubly instantiated, nodes sustain two separate oscillation frequencies and this may sometimes lead to uneven lags between successive spikes: (compare the double instantiation of “Man” with that of “Lover” in Figures 2). This phenomenon is similar to bifurcation by *period doubling* (Canavier, Clark, & Byrne, 1990). A stable oscillatory state can lose its stability, giving rise to a new stable state with doubled period. This phenomenon, when repeated, often leads to chaos.

The solution in the above example is restricted to double instantiation because multiple instantiation puts extra constraints on working memory. The above example requires 6 windows of synchrony, which fill working memory. When multiple instantiation is needed INFERNET must split role nodes into different phases. How to take greater number of instantiations into account, is currently being studied.

Performance of INFERNET

Figure 5 shows the performance of the computational implementation of INFERNET. The task was to find “Whose love is reciprocated?” and “Who is jealous of whom?” when given “John loves Louise and Louise loves John” and “John loves Louise and Louise loves Gray.” The performance of the system is measured by the percentage of correct responses and by the time taken by the system to set the correct bindings.

Two variables were manipulated: the amount of distribution (Overlapping vs. Non overlapping distributions) and the presence or absence of noise in the system. In this experiment, each concept is composed of 16 nodes. In the Non-overlapping condition, no concept shares nodes with other concepts, whereas in the overlapping condition, each concept shares 4 nodes with two other concepts which will never be bound to the same role. Noise is added at each time step. Experiments consisted of 20 trials for each of the four conditions. Figure 5 shows that, in general, overlapping distributions reduce the percentage of correct answers, and when the response is correct, response time decreases (if there is no noise). In the task tested, most concepts are doubly instantiated and the representational overlap means that more instantiations will occur.

Consequently, certain nodes must be assigned to more than two window of synchrony. For example, "Louise" nodes must be synchronized with "Lover" and "Lovee" nodes, but if "Louise" shares nodes with "Love", these shared nodes must fire in additional windows of synchrony. Since these nodes cannot oscillate faster than 100Hz, some of the required spikes cannot occur and the proportion of correct answers thus decreases.

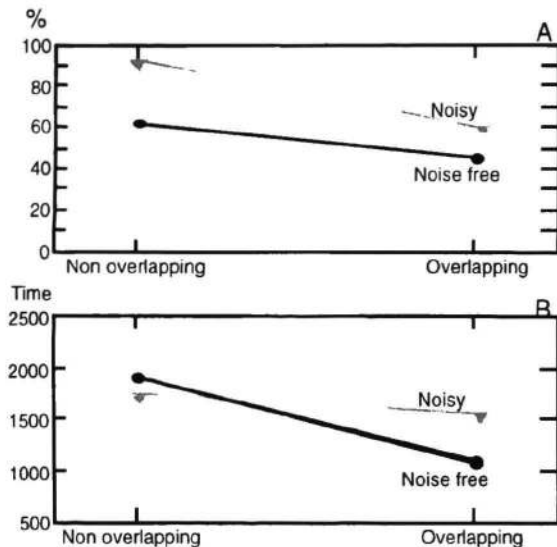


Figure 5: INFERNET results for a task requiring multiple instantiation. (A) shows the percentage of correct responses, (B) displays the time taken by the system to learn the correct bindings.

Why does binding fixation convergence time decrease in the noise-free overlapping condition? Postsynaptic nodes require the conjunction of activation at a precise time to fire. If the conjunction involves input from different concepts and if these concepts share nodes in common, the increase in firing rate increases the chance of having a conjunction of activation that causes the firing of the postsynaptic node. On the other hand, this will also increase the number of inappropriate firings of these postsynaptic nodes. This increases response errors and the system rapidly reaches a local minimum. When noise is added, it provides a means of escaping from these local minima, thus improving the frequency of correct responses, but also increasing response times. Noise makes the system more erratic before reaching a stable point. It allows exploration of a larger part of the space, which takes time, but also improves the chance of finding the best answer. The general effect of noise is similar to the phenomenon of stochastic resonance (see Levin & Miller 1996).

Conclusions

Multiple instantiation poses problems not only for connectionism, but for humans as well. However, with adequate time, humans can represent data that involve multiple instantiation. In this paper, various solutions to the problem of multiple instantiation are discussed.

The first uses separate LTM and STM stores. Even if STM has a limited capacity, one can fill it with instances of the same concept. This solution does not predict people difficulties with multiple instantiation.

The second solution makes multiple copies of the same content in LTM. It increases the storage capacity requirement of LTM. It predicts an increase in reaction time proportional to the number of concept instances, and an abrupt disruption of response quality if the number of instantiation required by the task is superior to the number of available copies.

The third solution uses multiple oscillation frequencies which increases the load of STM. The number of instantiations is limited by STM capacity and by the range of possible oscillation frequencies. This solution predicts a decrease in response quality proportional to the number of instances required. This decrease is associated with increased STM load.

The last solution better reflects people's difficulties with multiple instantiation, although additional psychological studies are still needed.

Acknowledgments

This research was supported by the Belgian PAI Grant p4/19 Special thanks to Robert French and Elizabeth Thomas for their assistance in the work presented here.

References

- Abeles, M., Prut, Y., Bergman, H., Vaadia, E. & Aertsen, A. (1993). Integration, Synchronicity and Periodicity. In A. Aertsen (Ed.) *Brain Theory: Spatio-Temporal Aspects of Brain Function*. Amsterdam: Elsevier.
- Baddeley, A. D. (1966). Short term memory for word sequences as a function of acoustic, semantic and formal similarity. *Quarterly Journal of Experimental Psychology*, 18, 362-365.
- Barnden, J. A. (1994). On the connectionist implementation of analogy and working memory matching. In J. A. Barnden, & K. J. Holyoak (Eds.) *Advances in Connectionist and Neural Computation Theory Volume 3: Analogy, metaphor, and reminding*. Norwood, NJ: Ablex Publishing Corporation.
- Bookman, L. A. and Alterman, R. (1991). Schema recognition for text understanding: An analog semantic feature approach. In J. A. Barnden, & J. B. Pollack, (Eds.) *Advances in Connectionist and Neural Computation Theory Volume 1: High-level connectionist models*. Norwood, NJ: Ablex Publishing Corporation.
- Canavier, C. C., Clark, J. W. and Byrne, J. H. (1990). Routes to chaos in a model of a bursting neuron, *Biophys. J.*, 57, 1245-1251.
- Carreiras, M. and Santamaria, C. (1997). Reasoning about Relations: Spatial and Nonspatial Problems. *Thinking and Reasoning*, 3, 191-208.
- Cowan, N. (1998). Visual and auditory working memory capacity. *Trends in Cognitive Sciences*, 2, 77-78.
- Dyer, M. G. (1991). Symbolic NeuroEngineering for Natural Language Processing: A Multilevel Research

- Approach. In J. A. Barnden, & J. B. Pollack, (Eds.) *Advances in Connectionist and Neural Computation Theory Volume 1: High-level connectionist models*. Norwood, NJ: Ablex Publishing Corporation.
- Engel, A. K., Kreiter, A. K., König, P., & Singer, W. (1991). Synchronisation of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proc. Natl. Acad. Sci. U. S. A.*, 88, 6048-6052.
- Feldman, J. A. (1982). Dynamic Connections in Neural Networks. *Biol. Cybern.*, 46, 27-39.
- French, R. M. (1995). *The Subtlety of Sameness: A theory and computer model of analogy-making*. Cambridge, MA: MIT Press.
- Hawkins, R.D, Kandel, E. R. and Siegelbaum, S. A. (1993). Learning to modulate transmitter release: Themes and variations in synaptic plasticity. *Annu. Rev. Neurosci.*, 16, 625-665.
- Henderson, J. (1996). A Connectionist Architecture of Inherent Systematicity. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. (pp. 574-579) Mahwah, NJ: Lawrence Erlbaum Associates.
- Hummel, J. E. & Holyoak, K. J. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427-466.
- Kanwisher, N. G. (1987). Repetition Blindness: Type recognition without token individuation. *Cognition*, 27, 117-143.
- Konnerth, A., Tsien, R.Y., Mikoshiba, K. and Altman, J. (1996). *Coincidence detection in the nervous system*. Strasbourg: Human Frontier Science Program.
- Lange, T. E. (1992). Lexical and pragmatic disambiguation and re-interpretation in connectionist networks. *International Journal of Man-Machine Studies*, 36, 191-220.
- Lange, T. E. and Dyer, M. G. (1989). High-level inferencing in a connectionist network. *Connection Science*, 1, 181-217.
- Levin, J. E. and Miller, J. P. (1996). Broadband neural encoding in the cricket cercal sensory system enhanced by stochastic resonance. *Nature*, 380, 165-168.
- Lisman, J. E., & Idiart, M. A. P. (1995). Storage of 7 ± 2 Short-Term Memories in Oscillatory Subcycles. *Science*, 267, 1512-1515.
- Mani, D. R. and Shastri, L. (1993). Reflexive Reasoning with Multiple Instantiation in a Connectionist Reasoning System with a Type Hierarchy. *Connection Science*, 5, 205-242.
- Morris, A. L. and Harris, C. L. (1997). A sublexical locus for repetition blindness: Evidence from illusory words. *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. (pp. 519-524) Mahwah: Lawrence Erlbaum Associates.
- Norman, D. A. (1986). Reflections on Cognition and Parallel Distributed Processing. In J. L. McClelland, & D. E. Rumelhart (Eds.) *Parallel Distributed Processing Vol. 2* Cambridge, MA: MIT Press.
- Roelfsema, P. R, Engel, A. K., König, P. & Singer, W. (1996). The role of neuronal synchronization in response selection: A biologically plausible theory of structured representations in the visual cortex. *Journal of Cognitive Neuroscience*, 8, 603-625.
- Shastri, L. and Ajjanagadde, V. (1993). From Simple Associations to Systematic Reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Science*, 16, 417-494.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annu. Rev. Physiol.*, 55, 349-74.
- Sougné, J. (1996). A Connectionist Model of Reflective Reasoning Using Temporal Properties of Node Firing. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*. (pp. 666-671) Mahwah, NJ: Lawrence Erlbaum Associates.
- Sougné, J. (1998). Connectionism and the problem of multiple instantiation. *Trends in Cognitive Sciences*, 2, 183-189.
- Sougné, J. and French, R. M. (1997). A Neurobiologically Inspired Model of Working Memory Based on Neuronal Synchrony and Rhythmicity. In J. A. Bullinaria, D. W. Glasspool, and G. Houghton (Eds.) *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representations*. London: Springer-Verlag.
- Wang, X. & Rinzal, J. (1995). Oscillatory and Bursting Properties of Neurons. In A. Arbib (Ed.) *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press.
- Wilson, M. & Shepherd, G. M. (1995). Olfactory Cortex. In A. Arbib (Ed.) *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press.