Computer Science Department Technical Report
University of California
Los Angeles, CA 90024-1596

# ARTIFICIAL NEURAL OSCILLATORS FOR INFERENCING

Trent E. Lange
Jacques J. Vidal
Michael G. Dyer

November 1989
CSD-890062

# Artificial Neural Oscillators for Inferencing†*

Trent E. Lange
Jacques J. Vidal
Michael G. Dyer

Computer Science Department
University of California, Los Angeles
Los Angeles, CA 90024, USA

## ABSTRACT

Connectionist networks have been unable to perform high-level conceptual tasks because of their inability to handle the *variable binding* and *inferencing* problems. This paper proposes the use of artificial neural oscillators in a localist network to approach these problems. In this model, groups of relaxation oscillators with unique patterns of natural oscillation *frequencies* serve as signatures to identify the concepts bound to an oscillator "variable". Inferences are made as the frequency signatures representing variable bindings propagate across chains of phase-locking oscillators.

## 1. INTRODUCTION

Understanding natural language is a difficult task, often requiring a reader to make multiple inferences to understand the motives of actors and to connect actions that are unrelated on the basis of surface semantics alone. An example of this is the sentence:

> *"John put the pot inside the dishwasher because the police were coming."* (S1)

A complex plan/goal analysis of S1 must be made to understand the actors' actions and disambiguate *"pot"* to Marijuana by overriding the local context of *"dishwasher"*. This requires the ability to represent dynamic role-bindings and quickly propagate them for inferencing while combining contextual evidence to perform the disambiguation.

### 1.1. Distributed Spreading-Activation Networks

Distributed connectionist models, such as those of [McClelland & Kawamoto, 1986] and [Touretzky & Hinton, 1988], have lately been receiving much interest, in part because their massively parallel networks of simple processing elements suggest correspondence with the way information is processed in the brain. Despite this attention, no distributed network model has yet exhibited the ability to handle natural language input having complexity even near to that of S1. The primary reason for this current lack of success is the inability to perform dynamic role-bindings and to propagate these binding constraints during inferencing. Distributed networks, furthermore, remain sequential at the knowledge level (i.e. can select and fire only one rule at a time) [Dyer, 1989] and lack the representation of structure needed to handle complex conceptual relationships [Feldman, 1989].

### 1.2. Localist Spreading-Activation Networks

Localist spreading-activation models, such as those of [Cottrell & Small, 1982], [Waltz & Pollack, 1985], and [Shastri, 1988], also use massively parallel networks of simple processing units. Knowledge is represented in localist networks by simple computational nodes and their interconnections, with each node standing for a distinct concept. Activation on a conceptual node represents the amount of *evidence* available for that concept in the current context.

Unlike distributed networks, localist networks are parallel at the knowledge level and are able to represent structural relationships between concepts. Because of this, many different inference paths can be pursued simultaneously; a necessity if the speed of language understanding exhibited by people is to be accounted for.

Unfortunately, however, the evidential activation on the conceptual nodes of previous localist networks gives no clue as to *where* that evidence came from. Because of this, previous localist models have had no more success than distributed models at handling dynamic, non-local bindings — and thus remain unsuited to higher-level knowledge tasks where inferencing is required.

## 2. A VALUE-PASSING MODEL: ROBIN

ROBIN (ROle Binding and Inferencing Network), introduced in [Lange & Dyer, 1988, 1989a] and explained in detail in [Lange & Dyer, 1989b], is our localist spreading-activation model that retains the advantages of previous localist approaches but, in addition, handles dynamic role-binding for inferencing. Nodes in ROBIN's network are simple computational elements, whose activation represents one of two things:

*Evidential activation* — Activation that indicates the likelihood that a concept is selected in the current context.

*Signature activation* — Uniquely-identifying activation that represents and allows the propagation of dynamic role-bindings for inferencing.

ROBIN uses structured connections of nodes using both evidential and signature activation to encode frames [Minsky, 1975]. Each frame (schema) has one or more role nodes, with each role having expectations and logical constraints on its fillers. Every frame can be related to one or more other frames, with pathways between corresponding roles for inferencing. Evidential and signature activation spread from frame to related

---

†To appear in the *Proceedings of the International Workshop on Neurocomputers and Attention*, Moscow, USSR, September 18-22, 1989.

*The first and third authors have been supported in part by a contract with the JTF program of the DOD and grants from the W.M. Keck and ITA Foundations. The second author has been supported in part by grants from Aerojet Electro-Systems and NASA. The model has been implemented on an Apollo DN4000, donated by an Apollo Computer Inc. equipment grant to the UCLA interdisciplinary Cognitive Science Research Program.

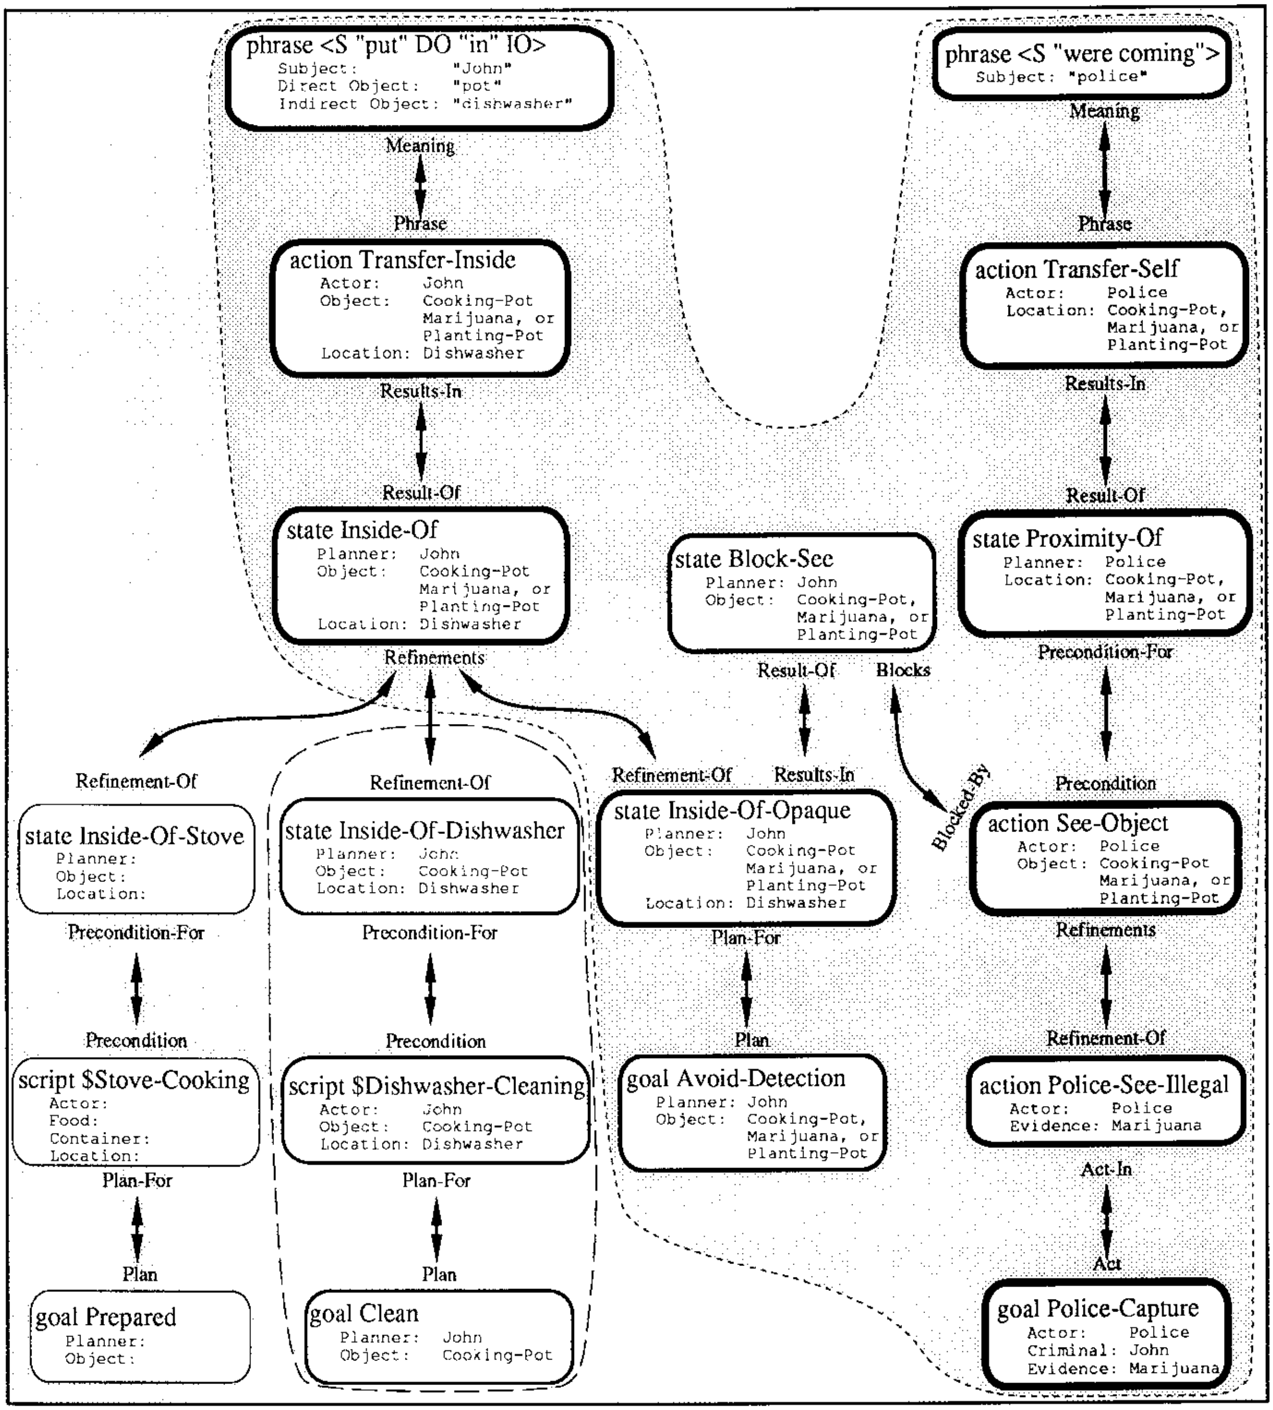


**Figure 1.** Overview of a ROBIN semantic network showing inferences dynamically made after S1 has been presented. Thickness of frame boundaries shows the amount of evidential activation for the frames. Bindings of frame roles have been realized through propagation of signature activation. Darkly shaded area indicates the most highly-activated path of frames representing the most probable plan/goal analysis of the input. Dashed area shows the discarded dishwasher-cleaning interpretation. Frames outside of both areas show a very small portion of the network that received no evidential or signature activation from either phrase.

frame when the constraints on their role fillers are met, thus automatically instantiating other frames and performing the processes of inferencing and frame selection.

For instance, in order to understand the sentence *"John put the pot inside the dishwasher because the police were coming."*, ROBIN makes the following inferences:

I1: If the police see John's marijuana, then they will know that he possesses an illegal object (since marijuana is an illegal substance).

I2: If the police know that John is in possession of an illegal object, then they will arrest him, since possessing an illegal object is a crime.

I3: John does not want to get arrested.

I4: John has the goal of stopping the police from seeing his marijuana.

I5: The police coming results in them being near John and his marijuana.

I6: The police being near John's marijuana enables them to see it.

I7: John's putting the marijuana inside the dishwasher results in the marijuana being inside the dishwasher.

I8: The marijuana is inside an opaque object (the dishwasher).

I9: Since the marijuana is inside an opaque object, the police cannot see it, thus satisfying John's goal.

Figure 1 shows a segment of the semantic network embedded in ROBIN after input for sentence S1 has been presented. The network has made inferences I1-I9, with most being shown in the figure. For example, I8 (the inference that the Marijuana is inside of an opaque object) is represented by the instantiation of state Inside-Of-Opaque. The role-bindings of the frames shown were instantiated dynamically with signature activation, with the final interpretation selected being the most highly-activated evidential path of frames inside the darkly shaded area.

### 2.1. The Structure of ROBIN

The processing units in ROBIN's networks each perform a simple computation on their inputs: summation, summation with thresholding and decay, or maximization. The connections between units are weighted and either excitatory or inhibitory.

As in previous localist models, ROBIN's networks have a node for every known concept in the network. Relations between concepts are represented by connections between their respective nodes, with the activation of a node corresponding to the amount of evidence (evidential activation) that exists in the current context for that concept. ROBIN's networks, however, have additional structure to represent dynamic role-bindings and to handle inferencing.

### 2.2. Signature Activation In ROBIN

As mentioned earlier, every conceptual node in ROBIN's localist network has associated with it an identification node broadcasting a stable, uniquely-identifying activation pattern, called its *signature*. A dynamic binding is created when a role's *binding node* has an activation that matches the activation of the bound concept's signature node.

For instance, in Figure 2, the virtual binding of the Actor role node of action Transfer-Inside to John is represented by the fact that its binding node, the solid black circle, has the same activation (3.1) as John's signature node.

### 2.3. Propagation of Signatures For Inferencing

The most important feature of ROBIN's signature activations is that the model passes them, as activation, across sometimes long paths of binding nodes to handle the non-local role-bindings necessary for inferencing. Figure 3 illustrates how the structure of the network automatically accomplishes this in a ROBIN network segment that implements a portion of the semantic network of Figure 1.

In Figure 3, the virtual binding of the Object role of frame Transfer-Inside to objects Marijuana and Cooking-Pot (from the phrase *"John put the pot inside the dishwasher"*) is represented by the fact that its corresponding binding nodes have the same activations (6.8 and 9.2) as the objects' signatures.

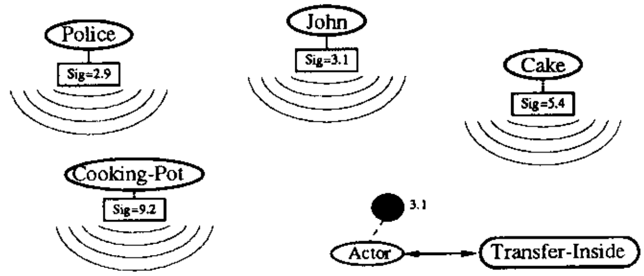


**Figure 2.** Several concepts (ovals) and their uniquely-identifying signature nodes (rectangles) are shown, along with the Actor role of the Transfer-Inside frame. The Actor role has a *virtual binding* to John because its binding node (black circle) has the same activation (3.1) as John's signature.

Each of the binding nodes calculates its activation as the maximum of its inputs, thus allowing the signature activations to be propagated and preserved across the paths between the corresponding roles of related frames. For example, in Figure 3, the signature activations have reached the binding nodes of Inside-Opaque's Object, showing that the network has inferred that either Marijuana or Cooking-Pot might be inside of something that is opaque (the dishwasher). The greater evidential activation of Marijuana (thicker oval) reveals the currently preferred interpretation of *"pot"*.

### 2.4. Selection of Ambiguous Role-Bindings

Note that *all* ambiguous meanings of a word are bound to a role with signature activation (Figures 1 and 3). The network's interpretation of which binding is selected at any given time is the binding whose concept has greater evidential activation. Because all candidate bindings propagate in the network, with none being discarded until processing is completed, ROBIN is able to handle meaning re-interpretations without resorting to backtracking. For further details on the use of signature activation for role-bindings see [Lange & Dyer, 1989b].

During the interpretation of S1, Cooking-Pot initially receives more evidential activation than Marijuana by connections from the highly stereotypical usage of the dishwasher for the Clean goal. The network's decision between the two candidate bindings at that point would be that it was a Cooking-Pot that was Inside-Of the Dishwasher. However, reinforcement and feedback from the inference paths generated by the Police's Transfer-Self eventually cause Marijuana to win out. The final selection of Marijuana over the Cooking-Pot bindings is represented simply by the fact that Marijuana has greater evidential activation. The resulting most highly-activated path of nodes and non-local bindings represents the plan/goal analysis in Figure 1. A more detailed description of ROBIN's network structure and capabilities can be found in [Lange & Dyer, 1989b].

## 3. SYNCHRONIZABLE OSCILLATORS

The vast majority of connectionist models employ value-passing processing elements that calculate a binary or continuous activation based upon a function of their inputs. Their activation is passed through an output function to produce a numeric output which becomes the input to other nodes. These value-passing elements and their connections have been "neurally-inspired" [McClelland & Rumelhart, 1986], with their numeric outputs likened to a simple "frequency-coded" representation of real neurons' action potential activity.

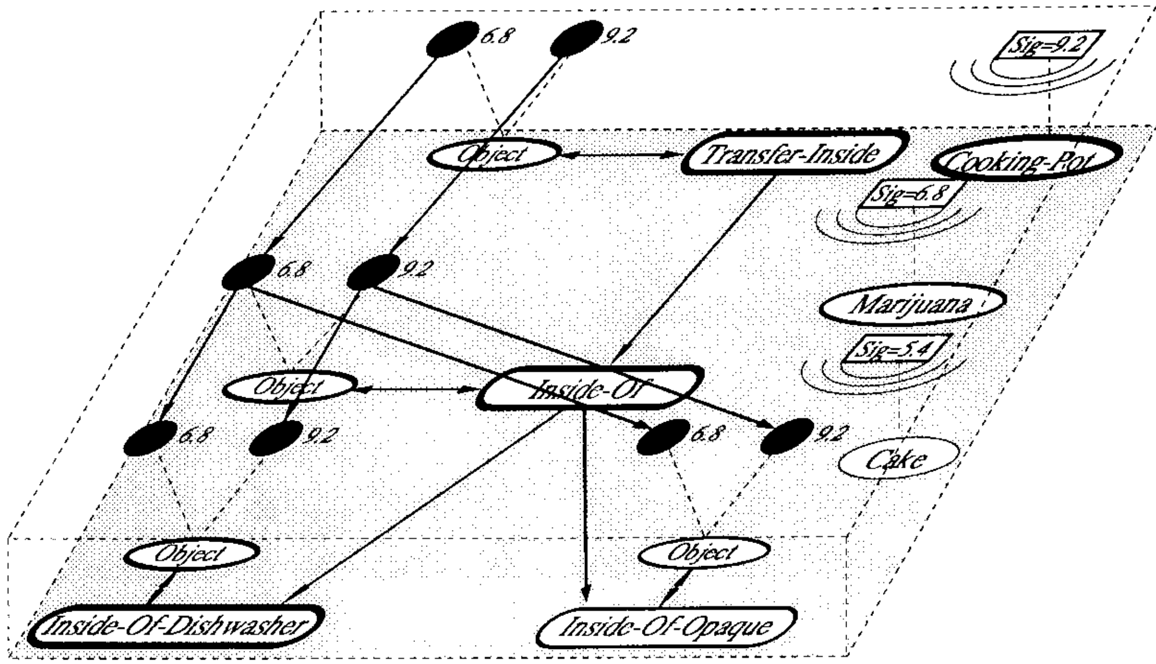


**Figure 3.** Simplified ROBIN network segment showing parallel paths over which evidential activation (bottom plane) and signature activation (top plane) are spread for inferencing. Signature nodes (rectangles) and binding nodes (solid black circles) are in the top plane. Thickness of conceptual node boundaries (ovals) represents their level of evidential activation after quiescence has been reached for sentence S1. (The names on the nodes are not used by ROBIN in any way, being used simply to set up the network's structure initially and to aid in analysis.)

### 3.1. Neurophysiological Background and Inspiration

While the neurophysiological plausibility of connectionist models is open to debate, it may be pointed out that many of the valuable timing dynamics of real neurons are lost when the popular frequency-coded simplification is used. The importance of precise action potential timing and phase relationships between neighboring neurons has been illustrated in several studies.

Of particular interest here are the dynamics of groups of neurons that act as neural oscillators, generating stable patterns of spike trains. An example of this has been shown by [Segundo *et al.*, 1964]:

> "The mechanism described determines stable patterns in which, over a clearly defined frequency range, the output discharge is locked in phase and frequency..." [Segundo *et al.*, 1964]

The interactions of neighboring neural oscillators are particularly striking. The output of an oscillator receiving a strong input, either excitatory or inhibitory, from a neighboring oscillator will often adjust its timing until its output spikes are locked in both phase and frequency with its input [Segundo & Kohn, 1981]. This kind of synchronization is lost in the simple value-passing elements of traditional artificial neural networks.

Our new model employs relaxation oscillators that approximate some of the complex timing dynamics of their neural counterparts. We do not make any claims about the neurophysiological plausibility of our model, but do utilize some of the gross features of real neural oscillators as an improved mechanism for dynamic role-binding and inferencing.

### 3.2. A Model of a Relaxation Oscillator

Figure 4 shows a simple model of a relaxation oscillator of [Vidal & Haggerty, 1987]. In a coarse approximation to similar mechanisms in real neurons, the relaxation oscillator accumulates activation energy until a threshold ($\mathbf{E_C}$) is reached. At that point the stored energy is abruptly released, accompanied by a brief output spike.

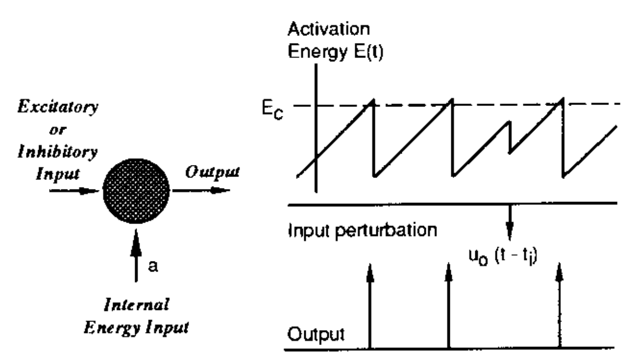


**Figure 4.** Relaxation Oscillator with perturbation input.

As shown in Figure 4, each relaxation oscillator possesses its own constant rate of energy influx, **a**. In absence of any external inputs, the oscillator will fire with the natural period:

$$\mathbf{T} = \frac{\mathbf{E_c}}{\mathbf{a}}$$

Between spikes, when external input in the form of input pulses at time $t_j$ arrives, the activation energy of the oscillator is calculated by:

$$E(t - t_0) = a(t - t_0) + \sum_{1}^{J} w_j \cdot u_0(t - t_j) \ ; \ E \leq E_c$$

where $t_0$ represents the instance of the last firing of the cell, $t_j$, ($j = 1,2,...J$), the instants of pulses arriving from other cells, $w_j$ the input weight for pulse j, and $u_0(t)$ represents the unit impulse at $t=0$.

### 3.3. Phase Locking

Two oscillators that are phase-locked, or synchronized, with each other, fire with the same period at a constant phase apart. [Vidal & Haggerty, 1987] showed how a master-slave relationship between two relaxation oscillators forces the slave oscillator to synchronize with its master. In Figure 5, master oscillator **A** (having natural period $T_A$) inhibits oscillator **B** (having natural period $T_B$). The energy increment or decrement introduced by the arrival of an input pulse is a function of its relative arrival time with respect to the last output firing. As a result of this, regular inhibitory spikes from A cause the phase of B to self-adjust until an equilibrium is reached with new period $T_A$. The precise value of the inhibitory connection's weight affects only the phase with which B's output spikes follow A's, and not its ability to take on A's firing frequency.

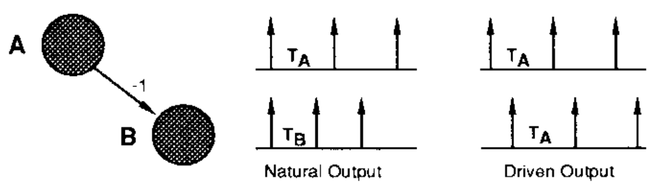


**Figure 5.** When one oscillator (A) drives another oscillator (B), the driven oscillator will tend to phase-lock with the first by taking on its natural period. Each upward-pointing arrow represents a single output spike.

In general, the synchronization dynamics present an attractor for each rational frequency pair. This type of multiple phase-locking has been demonstrated in living invertebrate neurons [Segundo & Kohn, 1981], as well as in simulations.

## 4. SYNCHRONIZABLE OSCILLATORS FOR ROLE-BINDING

The two critical features of ROBIN's signature activations are that they serve as stable patterns to uniquely identify concepts and that they can be propagated across long paths of binding nodes to allow dynamic role-binding and inferencing.

As previously described, ROBIN uses simple value-passing elements similar to most other artificial neural networks. A concept's signature is produced by a signature node that outputs a constant numeric value that is unique to that signature. To allow signatures to propagate, binding nodes are elements whose activation and output are equal to the maximum of their inputs, with each input connection having unit weight. Signature activations are thus preserved as they are propagated across paths of binding nodes for inferencing (as in Figure 3).

We can easily envision how the nodes over which evidential activation pass could be "implemented" in neural-like hardware. Since their activation serves only as an approximate measure of how likely a concept is in the forefront of memory in a given context, it is not necessary for them to exactly calculate the sum of their inputs. A simple integration of their input evidence will easily suffice, with levels of activation represented by frequencies of output firing.

As we examine the model further, however, the binding nodes that exactly calculate the maximum of their inputs appear more nettlesome. It is critical that these nodes calculate the maximum of their inputs exactly, and that their inputs all have precisely unit weight. A small deviation from either unit input weight or perfect maximum calculation will cause different levels of activation; therefore possibly representing a different signature and hence role-binding. The accumulative effect of such noise as signatures are passed along long chains of binding nodes could be devastating to the inferencing process, even if there is a margin of error between signatures. Small errors in signature propagation would become less troublesome if signatures were distributed patterns of activation rather than single numbers, as we proposed in [Lange & Dyer, 1988, 1989b], but the problem would still remain.

### 4.1. Relaxation Oscillators for Signatures and Bindings

We have developed an extension of ROBIN whose elements communicate via action potentials rather than the simple values characteristic of most connectionist models. Synchro-ROBIN uses relaxation oscillators to both produce its concept-identifying signatures and propagate them across long paths of nodes for dynamic role-binding and inferencing. The complex dynamics of action potential timing actually increases the model's resistance to noise.

Synchro-ROBIN uses a separate relaxation oscillator for each signature-producing node. Each *signature oscillator* has a characteristic level of internal energy influx or firing threshold so that the natural frequency at which it produces output spikes is different than every other oscillator signature. The signature oscillators are cut off from external input during the short-term inferencing process so that they may produce output spikes at their natural frequencies.

Binding nodes in Synchro-ROBIN are also relaxation oscillators. A virtual dynamic binding is made when a *binding oscillator* is locked in phase with an oscillator signature.

In Figure 6, the virtual binding of the Actor role of action Transfer-Inside is represented by the fact that output spikes of its binding oscillator are temporarily locked in phase with the signature oscillator of John.

As in ROBIN, the oscillator signatures of Synchro-ROBIN are passed across long paths of interacting binding oscillators to handle inferencing. In Figure 7, the left binding oscillator of Transfer-Inside's Object role is locked in phase with the signa-

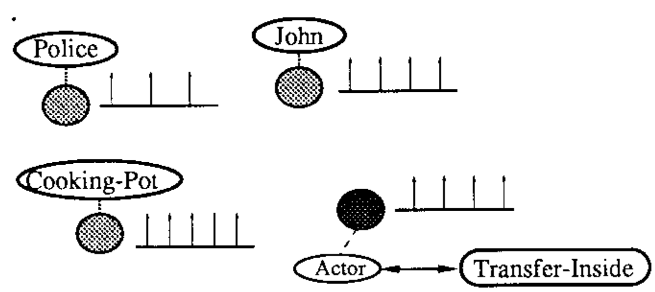


**Figure 6.** Several concepts and their uniquely-identifying signature oscillator (light grey circles). The binding oscillator (dark grey circle) of the Actor role is a relaxation oscillator locked in phase with the signature frequency of John, thus representing the binding.

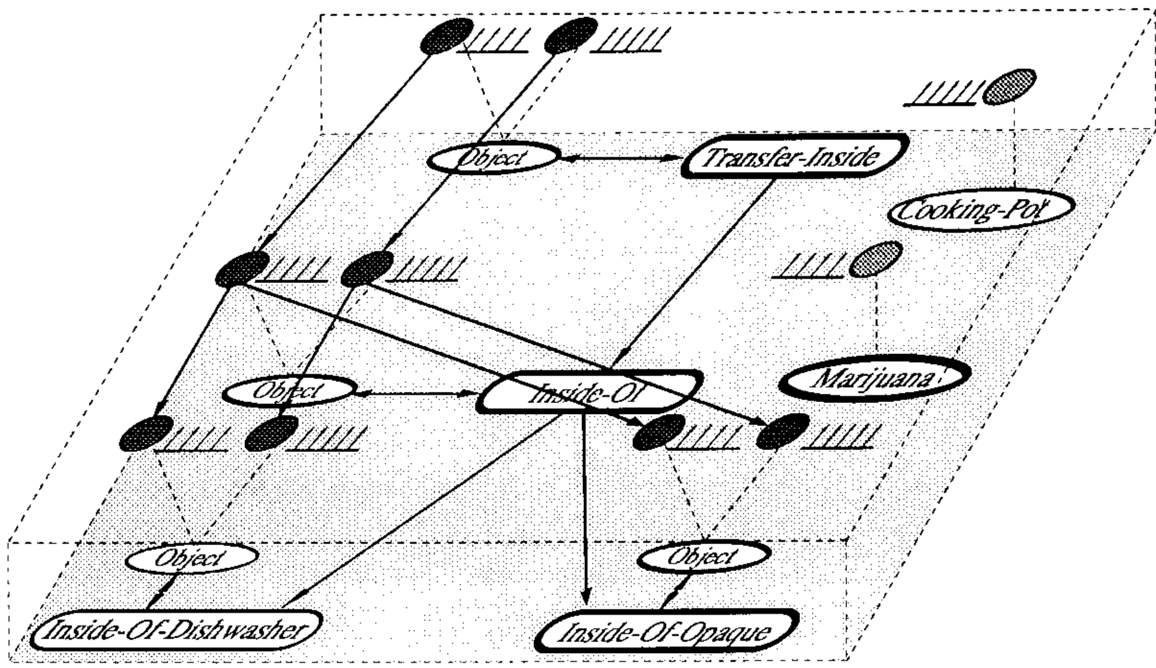


**Figure 7.** Simplified Synchro-ROBIN network segment shows parallel paths over which evidential activation and oscillator signature activation are spread for inferencing. Signature oscillators (light grey circles) and binding oscillators (dark grey circles) are in the top plane. The rest of the figure is the same as Figure 3, with outputs after quiescence has been reached in processing of sentence S1.

ture oscillator of Marijuana, representing the virtual binding from the phrase *"John put the pot inside the dishwasher"*. Once it is locked in phase, an inhibitory connection from it to the related left binding oscillator of Inside-Of's Object role causes that oscillator to lock in phase, making the inference that the Marijuana may be Inside-Of something. Similarly, the competing oscillator signature bindings of Cooking-Pot are spread across the parallel paths of the right binding oscillators. The oscillator signatures continue to propagate, through synchronization of corresponding binding oscillators, until stability is reached and the inferences of Figure 1 have been made.

### 4.2. A More Detailed Look

There are a couple of special concerns when using relaxation oscillators as signature frequencies for role-binding. First of all, as shown by [Vidal & Haggerty, 1987], the dynamics of oscillator synchronization presents an attractor for each rationally-valued frequency pair, forming a "devil's staircase" of stability regions. These stability regions become smaller as the ratios become larger or smaller than unity. Since we are interested in preserving the signature frequencies, the natural frequencies of the signature oscillators and the binding oscillators must be in a close enough range that their frequencies are locked in phase on a 1:1 basis.

A difficult problem is caused by the nature of inferencing. In general, the filler of a role can be inferred in several ways. The fact that a Cooking-Pot is inside of a dishwasher, for example, can be inferred from either the knowledge that somebody was putting it inside the dishwasher or that the $Dishwasher-Cleaning script was being used to clean it. Figure 8 illustrates the problem of two binding oscillators (A and B) feeding in as possible inference paths for another (C).

The problem is that the dynamics of oscillator attraction will often fail when two oscillators feed into another, even if both driving oscillators are themselves locked in phase. The extra inhibition may cause the receiving oscillator to either lock into an output frequency slower than both inputs or to reach no stable frequency at all.

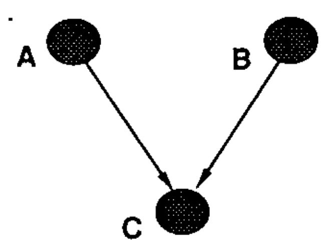


**Figure 8.** The binding of binding oscillator C can be inferred from either binding A or binding B.

In natural language understanding, however, it is generally observed that one frame will not be inferred from two others at exactly the same instance in time. In sentence S1, the fact that the Marijuana was Inside-Of the dishwasher was first inferred from the fact that it was Transfer-Inside of it. What is necessary, then, is to shut off the alternative inference paths once the inference (phase-locked signature oscillation) has been made.

Figure 9 shows the excitatory and inhibitory connections between the binding oscillators of Figure 8 to guarantee a single phase-locking at once. Synchro-ROBIN's binding oscillators are modified relaxation oscillators that do not have their own internal energy input. They are therefore naturally quiescent when they are not bound to a signature frequency. Instead, they have a gated excitatory input (link 1) from an external oscillator (D) that is naturally firing at a very rapid rate. The weighted excitatory synapse from this oscillator, when active, acts as the internal energy input for the binding oscillator.

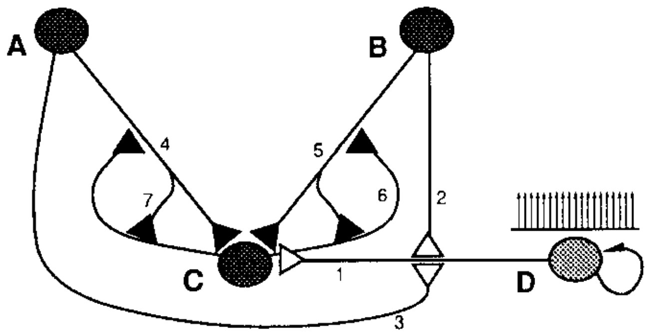


**Figure 9.** Interaction between the binding oscillators of Figure 8 used to assure gating of only one phase-locking at a time. Solid black triangular connections are strong inhibitory connections, while white triangular connections are excitatory.

When neither **A** nor **B** are active, **C** cannot be inferred and so should itself remain inactive. Since the gated external input link from **D** to **C** (link 1) can only be active when it receives additional excitatory input from **A** (link 2) or **B** (link 3), **C** will not receive its "internal energy input", and will indeed remain inactive and non-firing.

Consider now what happens when binding oscillator **A** becomes locked in phase (from some other inference), while **B** remains inactive. First, the gated connection from **D** to **C** (1) will become active, allowing **C** to start towards its natural frequency. At the same time, the inhibitory connection from **A** (4) will cause it to lock in phase with **A**'s signature frequency, performing the inference.

Binding oscillator **B** becoming active at a later time will have no effect on **C**, however, because its inhibitory link to **C** (5) has itself been disabled by the activity of **C** (6). The original phase-locking inhibition from **A** continues, however, since the inhibition of link 7 has stopped link 4 from being inhibited by the output spikes of **C**.

Note that there are no direct connections between mutually-exclusive binding oscillators. If such connections were employed, then $O(n^2)$ connections would be required for $n$ oscillators feeding into **C**. The scheme in Figure 9, however, requires only $O(n)$ inhibitory connections.

With these types of inhibitory gated interconnections, Synchro-ROBIN's binding oscillators receive input from and lock into phase with the first related binding oscillator that becomes active with an signature frequency, thus dynamically propagating role-bindings for inferencing while avoiding crosstalk. The precise timing of output firings, lost with standard value-passing connectionist models, is crucial to these abilities and to the model's imperviousness to imperfect connection weights.

## 5. CURRENT STATUS AND FUTURE WORK

ROBIN has been fully implemented using value-passing elements in the DESCARTES connectionist simulator [Lange *et al.*, 1989]. DESCARTES is a development environment written in the Common Lisp Object System that allows the flexible simulation of large-scale heterogeneous connectionist networks. ROBIN's inferencing, plan/goal analysis, schema instantiation, disambiguation, and re-interpretation abilities have been successfully tested in two small domains, using natural language inputs of one or two sentences in length that have been syntactically preprocessed.

The use of relaxation oscillators for signature and binding oscillators in Synchro-ROBIN has also been implemented in DESCARTES. Dynamic role-bindings are propagated for inferencing by the synchronization of long paths of binding oscillators using the gating structure of Figure 9.

In summary, there are several directions for future research:

*Exploring More of the Features of Real Neural Oscillators:* The relaxation oscillators used in Synchro-ROBIN do not approach the richness of real neurons. More realistic models, such as those of [Segundo & Kohn, 1981] or [Kirillov *et al.*, 1988], might lead to more insight into how synchronization is performed.

*Distributing Evidential Conceptual Nodes:* The structure of the network over which evidential activation is passed still uses value-passing nodes. Replacing value-passing elements with groups of spike-firing units for each evidential node could help approach the "grandmother node" problems inherent to localist networks.

*Network structure acquisition:* non-local bindings allow Synchro-ROBIN to create novel network instances over its pre-existing structure. Over time, repeated instantiations should cause modification of weights and recruitment of underutilized units to alter network structure.

### 5.1. Distributed Signature Frequencies

There are limitations to using a single oscillator with a uniquely-identifying frequency to represent each signature. Large models could conceivably have thousands or hundreds of thousands of separate concepts that they could recognize (such as Marijuana, Cooking-Pot, Catfish, Guppy, John, John-Wayne, John-Kennedy, etc). If each concept is to have a single unique signature frequency, then there must be equally as many separate frequencies. The time to reliably transmit and lock into frequencies ranging up to the hundreds of thousands or millions would clearly rule out any models' pretenses to psychological plausibility.

A better solution is to use *groups* of oscillators to form distributed representations of signatures. Figure 10 shows an example where the signature for John is represented by the top group of signature oscillators and their natural oscillation frequencies. Similar concepts would have similar distributed patterns of oscillation frequencies, with a unique pattern for each

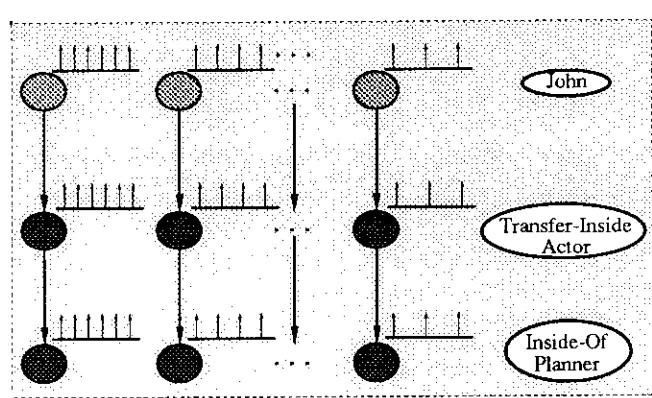


**Figure 10.** Groups of oscillators forming a distributed representation of John's signature. These distributed signature oscillations would also be propagated for inferencing through phase-locking of connected binding oscillators.

signature. A first pass at this might entail the use of microfeature-like patterns, as in the model of [McClelland & Kawamoto, 1986], but it would be preferable to have the signature patterns learned over time, as done by [Miikkulainen & Dyer, 1988] for value-passing backpropagation networks.

As shown in Figure 10, signatures represented with groups of relaxation oscillators would be propagated for inferencing through phase-locking of connected binding oscillators in exactly the same way as for single signature frequencies.

## 6. CONCLUSIONS

Our value-passing localist model, ROBIN, is a spreading-activation model that approaches many of the problems of natural language understanding, including those of inferencing and frame selection. The activation on the network's simple computational nodes is of one of two types: (a) *evidential activation*, to perform disambiguation by indicating the likelihood that a concept is selected, and (b) *signature activation*, to uniquely identify concepts and allow the representation and propagation of dynamic role-bindings not possible in previous connectionist models.

This paper describes how the signature activations of ROBIN can be handled with processing elements that communicate via output firings rather than the simple frequency-coded values of most connectionist models. Our new model, Synchro-ROBIN, uses relaxation oscillators to produce the signatures that identify concepts. A role-binding is created when a binding oscillator becomes locked in phase with the signature frequency of the bound concept. Chains of binding oscillators allow role-bindings, in the form of signature oscillation synchronization, to be propagated for inferencing in a manner that is more robust to noise than when using simple value-passing elements. Using parallel paths of relaxation oscillators to propagate signature frequencies and conventional value-passing nodes to handle evidential activation, Synchro-ROBIN is able to perform much the inferencing necessary in high-level conceptual tasks.

## REFERENCES


Cottrell, G. & Small, S. (1982): A Connectionist Scheme for Modeling Word-Sense Disambiguation. *Cognition and Brain Theory, Vol. 6*, p. 89-120.

Dyer, M. G. (1989): Symbolic Neuroengineering for Natural Language Processing: A Multi-Level Research Approach. In J. Barnden and J. Pollack (eds.), *Advances in Connectionist and Neural Computation Theory*, Ablex Publishing. In press.

Feldman, J. A. (1989): Neural Representation of Conceptual Knowledge. In L. Nadel, L. A. Cooper, P. Culicover, and R. M. Harnish (eds), *Neural Connections, Mental Computation*. Cambridge, MA: The MIT Press.

Kirillov, A.G., Borisyuk, G.N., Borisyuk, R.M., Kovalenk, Ye.I., Kryukov, V.I., Makarenko, V.I., & Chulaevsky, V.A. (1988): Short-Term Memory as a Metastable State. IV. A Model of Neural Oscillator for a Unified Submodule. In R. Trappl (Ed.), *Cybernetics and Systems '88*, Kluwer Academic Publishers, p. 991-998.

Lange, T. & Dyer, M. G. (1988): Dynamic, Non-Local Role-Bindings and Inferencing in a Localist Network for Natural Language Understanding. In David S. Touretzky (ed.), *Advances in Neural Information Processing Systems I*, p. 545-552, San Mateo, CA: Morgan Kaufmann (Collected papers of the IEEE Conference on Neural Information Processing Systems — Natural and Synthetic, Denver, CO, November 1988).

Lange, T. & Dyer, M. G. (1989a): Frame Selection in a Connectionist Model of High-Level Inferencing. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*, Ann Arbor, MI, August 1989.

Lange, T. & Dyer, M. G. (1989b): High-Level Inferencing in a Connectionist Network. *Connection Science, Vol. 1 (2)*. In press.

Lange, T., Hodges, J., Fuenmayor, M., & Belyaev, L. (1989): DESCARTES: Development Environment For Simulating Hybrid Connectionist Architectures. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*, Ann Arbor, MI, August 1989.

McClelland, J. L. & Kawamoto, A. H. (1986): Mechanisms of Sentence Processing: Assigning Roles to Constituents of Sentences. In McClelland & Rumelhart (eds.), *Parallel Distributed Processing: Vol 2*. Cambridge, MA: The MIT Press.

McClelland, J. L. & Rumelhart, D. E.. (1986): Biological Mechanisms. In Rumelhart & McClelland (eds.) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol 2*. Cambridge, MA: The MIT Press, p. 327-331.

Miikkulainen, R. & Dyer, M. G. (1988): Forming Global Representations with Extended Backpropagation. *Proceedings of the IEEE Second Annual Conference on Neural Networks (ICNN-88)*, San Diego, CA.

Minsky, M. (1975): A Framework for Representing Knowledge. In P.H. Winston (ed.), *The Psychology of Computer Vision*. New York: McGraw-Hill, p. 211-277.

Segundo, J. P., Perkel, D. H., Schulman, J. H., Bullock, T. H., & Moore, G. P (1964): Pacemaker Neurons: Effects of Regularly Spaced Synaptic Input. *Science*, vol. 145, p. 61-63.

Segundo, J. P. & Kohn, A. F. (1981): A Model of Excitatory Synaptic Interactions Between Pacemakers. Its Reality, its Generality, and the Principles Involved. *Biological Cybernetics*, vol. 40, p. 113-126.

Shastri, L. (1988): A Connectionist Approach to Knowledge Representation and Limited Inference. *Cognitive Science, Vol. 12 (3)*, p. 331-392.

Touretzky, D. S. & Hinton, G. E. (1988): A Distributed Connectionist Production System. *Cognitive Science, Vol. 12 (3)*, p. 423-466.

Vidal, J. & Haggerty, J. (1987): Synchronization in Neural Nets. *Proceedings of the IEEE Conference on Neural Information Processing Systems — Natural and Synthetic (NIPS-87)*, Denver, CO.

Waltz, D. & Pollack, J. (1985): Massively Parallel Parsing: A Strongly Interactive Model of Natural Language Interpretation. *Cognitive Science, Vol. 9 (1)*, p. 51-74.