

How to Solve the Problem of Compositionality by Oscillatory Networks

Markus Werning (markus.werning@uni-erfurt.de)

Department of Philosophy, P. O. Box 900221

D-99105 Erfurt, Germany

Abstract

Cognitive systems are regarded to be compositional: The semantic values of complex representations are determined by, and dependent on, the semantic values of primitive representations. Both classical and connectionist networks fail to model compositionality in a plausible way. The paper introduces oscillatory networks as a third alternative. It provides neurobiological evidence for the adequacy of those networks and argues that they are compositional. Oscillatory networks combine the virtues and avoid the shortcomings of classical and connectionist architectures.

Compositionality and Systematicity

Minds have the capacity to compose contents. Otherwise, they would not show a systematic correlation between representational capacities: If a mind is capable of certain intentional states in a certain intentional mode (perception, thought, imagination, preference, etc.), it most probably is also capable of other intentional states with related contents in the same mode. The capacity to see a red square in a green circle, e.g., is statistically highly correlated with the capacity to see a red circle in a green square. To explain this empirical phenomenon, which is closely related to the well-known binding problem, compositional operations are postulated. They enable the system to build complex representations from primitive ones so that the semantic value of the complex representation is determined by, and dependent on, the semantic values of the primitives. Several theories have been developed to meet the requirement of compositionality. Both classical and connectionist attempts suffer from severe deficits, though.

Fodor and Pylyshyn (1988) for one take recourse to a language of thought, which they link to the claim that the brain can be modeled by a Turing-style computer. A subject's having a mental representation, they believe, consists in the subject's bearing a computational relation to a mental sentence; it is a relation analogous to the relation a Turing machine's control head bears to the tape. Accordingly, the mind composes complex representations from primitive ones just the way a computer composes phrases from words: by concatenation. The trouble with classical computer models is well known and reaches from the frame problem, the problem of graceful degradation, and the problem of learning from examples (Horgan & Tienson, 1996) to prob-

lems that arise from the content sensitivity of logical reasoning (Gigerenzer & Hug, 1992).

To avoid the pitfalls of classicism, connectionist models have been developed. Some of them attempt to meet the compositionality constraint. Smolensky (1995) maps the terms and the syntax of a language homomorphically onto an algebra of vectors and tensor operations. Each primitive term of the language is assigned to a vector. Every vector renders a certain distribution of activity within the connectionist network. The syntactic operations of the language have tensor operations as counterparts. Barnden (1991) pursues a related approach. As far as syntax is concerned, some connectionist networks can completely implement compositional languages.

The kind of compositionality that is necessary for systematicity, however, focuses not on syntactic, but on semantic features. The capacity to think that a child with a red coat is distracted by an old herring is not correlated with the capacity to think that a child with an old coat is distracted by a red herring. The thoughts ought to be correlated, though, if syntactic composition was sufficient for systematicity. Although both thoughts are syntactically composed from exactly the same primitives by exactly the same operations, they are not correlated because *red herring* is idiomatic, i.e. because the mapping (*red, herring*) \rightarrow *red herring* is syntactically, but not semantically compositional. One may well have the capacity to think of red coats and old herrings even though one lacks the capacity to think of red herrings. We may infer, thus, that semantic compositionality is necessary for systematicity and that syntactic compositionality is not sufficient. The strategy to map the syntax of a systematic language homomorphically onto a connectionist network does not suffice to establish that the network itself is systematic.

To put the dilemma in a nutshell, connectionist models seem to be too weak to explain systematicity, whereas classical models are apparently too strong to be implemented by real brains. The rest of the paper will explore the option of something "in between" classical and connectionist architectures. The presented solution differs significantly from other approaches to the dilemma (Lange & Dyer, 1989; Shastri & Ajjanagadde, 1993; Hummel & Holyoak, 1997; Sougné 1999). Especially with respect to the representation of relations, the

presented model might have more plausible implications.

Constituency

A further argument provides us with a deeper insight into what's wrong with traditional connectionist networks and gives us a key how to match this deficit. Most semantic theories explain the semantic properties of internal representations either in terms of covariance, in terms of inferential relations, or in terms of associations. One may, e.g., hold that a certain internal state is a representation of redness because the state covaries with nearby instances of redness. This covariance relation is, of course, backed by the intrinsic and extrinsic causal properties of the *redness* representation. One may also hold that *bachelor* representations characteristically are such that the subject is disposed to infer *unmarried-man* representations from it. Those dispositions, again, are grounded in the causal properties of *bachelor* and *unmarried-man* representations. One may, thirdly, hold that the semantic value of the *cow* representation is determined by the fact that it is associated with representations like *milk*, *leather*, *mammal*, *grass* etc. The mechanism of association, too, supervenes on the causal properties of the representations in question. All of these theories have one principle in common: An internal representation has its semantic value because it has a certain causal role within the system (and – perhaps – the rest of the world). The question of how the semantic values of primitive representations determine the semantic value of complex representations, hence, leads to the question of how causal properties can be inherited from primitive to complex states. From chemistry we know that atoms determine the causal properties of molecules because atoms are *constituents* of molecules. Physics gives similar answers with regard to atoms and elementary particles. One can even make it a hard metaphysical point: If the causal properties of a state B are determined by, and dependent on, the causal properties of the states A_1, \dots, A_n and their relations to each other, then A_1, \dots, A_n are constituents of B. Here, constituents are conceived of as necessary parts: A is a constituent of B if and only if the following is necessary and generally true: If B occurs at a certain region of space at a certain time, then A occurs at the same region at the same time.

The failure of connectionist attempts, therefore, is that the homomorphism between language and network structure does not preserve the constituent relations within the language. The network counterparts of *brown* and *cow* aren't constituents of the network counterpart of *brown cow*. Since the homomorphism does not preserve constituent relations, it fails to transfer semantic compositionality: Although the operation $(brown, cow) \rightarrow brown\ cow$ is semantically compositional, the network operation $(h(brown), h(cow)) \rightarrow$




Figure 1: (a) Elementary oscillator consisting of a coupled pair of an excitatory (+) and an inhibitory unit (-) together with an input unit. t , time; $x(t)$, unit activity; $F(x)$ sigmoidal output function; w , coupling weight; τ , delay time; $i_e(t)$, external input. Subscripts: e, excitatory unit; i, inhibitory unit. (b) Oscillatory elements coupled by short-range synchronizing connections (dashed) and long-range desynchronizing connections (dotted), without interaction at crossings. The figure is meant to show the principle of coupled oscillators, rather than a particular connectivity pattern.

$h(brown\ cow)$ may not be semantically compositional (h being the homomorphism). If $h(brown)$ and $h(cow)$ aren't constituents of $h(brown\ cow)$ you cannot say: $h(brown\ cow)$ co-varies with brown cows *because* $h(brown)$ co-varies with brown things and $h(cow)$ co-varies with cows. If the constituent relations were preserved, you could say this. For the same reason, you are deprived of the possibility to explain the inferential and the associative properties of the complex representation on the basis of the inferential and the associative properties of the primitive representations.

Synchrony

Constituency is a synchronic relation, while causal connectedness is a diachronic relation. Whole and part co-exist in time, whereas causes and effects succeed in time. The reference to causal connections and the flow of activation within the network will, therefore, not suffice to establish constituent relations. What we, in addition, need is an adequate synchronic relation. Oscillatory networks provide a framework to define such a relation: the relation of synchrony between phases of activity. Synchrony and asynchrony are synchronic relations because the related phases of activity, coexist in time. An elementary oscillator is realized by coupling an excitatory unit with an inhibitory unit using delay connections. An additional unit allows for external input (figure 1a). Within the network, oscillatory elements are coupled by either short-range synchronizing connections or long-range desynchronizing connections




Figure 2: Scheme of a typical neuronal response aroused by a blue vertical, a red vertical, and a blue horizontal object. Circles with letters signify neurons with the property they indicate (V, H: horizontal, vertical; R, G, B, Y: red, green, blue, yellow). Like shadings signify synchronous activity. The phases of some *blue*-neurons are synchronous with the phases of some *vertical*-neurons (middle-shading), some phases of *vertical*-neurons are in synchrony with some phases of *red*-neurons (dark-shading), and some *blue*-neurons fire in phase with some *horizontal*-neurons (light-shading).

(figure 1b). A multitude of oscillators can be arranged in feature modules (e.g., the color module), employing appropriate patterns of connectivity. Given a certain selectivity of the input unit, each oscillator is designed to indicate a certain property (e.g., redness) within the feature domain. Oscillators for like properties are connected synchronizingly; those for unlike properties are connected desynchronizingly. The pattern of connectivity may as well reflect topographical features. The behavior of oscillatory networks have been studied in detail elsewhere (Schillen & König, 1994). Stimulated oscillatory networks, characteristically, show object-specific patterns of synchronized and desynchronized oscillators within and across feature modules. Oscillators that represent properties of the same object synchronize because oscillatory networks implement the Gestalt principles. Oscillators that represent properties of different objects desynchronize. We observe that for each represented object a certain phase of activity spreads through the networks. The phase pertains only to oscillators that represent the properties of the object in question. Assuming that elementary oscillators are models of neurons and that oscillatory networks are models of part of the visual cortex, the results of these studies support two hypotheses:

Indicativity. As part of the visual cortex, there are collections of neurons whose function it is to show activity only when an object in the perceptual field instantiates a certain property.

Synchrony. Neurons that belong to two collections indicative for the properties π_1 and π_2 , respectively, have the function to show activity synchronous with each other only if the properties π_1 and π_2 are instantiated by the same object in the perceptual field.

The hypotheses are supported by neurobiological evidence. The indicative function of neurons was discovered by Hubel and Wiesel (1962, 1968). Neurobiologists meanwhile have specified a great variety of

feature domains: color, orientation, direction of motion, speed, luminance, etc. Property indicative neuronal collections will, subsequently, be called π -collections, with π standing for the property the neurons of the collection indicate.

A number of experimental data support the hypothesis of synchrony (detailed review: Singer & Gray, 1995).¹ Synchrony of neurons (<200 μ m apart) within one column was recorded in many different species and cortical regions of awake and lightly anaesthetized animals, and can be observed in the local field potential as well as in the multi-unit and paired single-unit recordings (Gray & Singer, 1989; Kreiter & Singer, 1992). Intercolumnar synchrony of distant neurons (>2mm) was shown by simultaneously recording the activity of neurons in different parts of the cortex (Schwartz & Bolz, 1991). Its occurrence within and between visual areas depends upon whether the neurons are stimulated by single or separate objects. For example, synchrony is strong when two neurons in V1 with non-overlapping but collinear preferred orientations are stimulated by a single long bar moving across their receptive fields (Gray et al., 1989). It is weaker when they are stimulated by two short collinear bars moving in the same direction, and it is absent altogether when the two short bars move in opposite directions. These and other results support the view that the synchrony of distributed activity in the visual system implements the well-established Gestalt principles of perceptual grouping. The issue of object-binding as stated by the principle of synchrony is supported by evidence from the primary visual cortex of the cat (Engel, König, & Singer, 1991) and other animals. These experiments show that when two neurons with different orientation and direction preferences are stimulated by a single moving bar that is sub-optimal for both, then they synchronize, but when they are stimulated by two separate bars, each being optimal for one of the neurons, then they do not. The representational function of synchrony is supported by studies of binocular rivalry with awake strabismic cats (Fries et al., 1997). There has long been anatomical evidence for long-range horizontal connections in V1 (Rockland & Lund, 1983). Lowel and Singer (1992) observed that these connections play a synchronizing role. Figure 2 provides a schematic overview.

Algebra

Oscillatory networks that implement the two hypotheses can be given an abstract algebraic description:

$$\mathbf{N} = \langle N_i, N_p, N_s; \phi_1, \dots, \phi_m; F_1, \dots, F_n; \approx, \nabla, \epsilon, \wedge \rangle.$$

¹ O'Keefe and Recce (1993), Wehr and Laurent (1996), Gawne, Kjaer and Richmond (1996) assume a more critical attitude with respect to the role of synchrony in object binding.

Below, this algebra will be shown to be isomorphic to a systematic language. The primitive entities of the algebra are (i) the phases of activity picked out by the symbols $\varphi_1, \dots, \varphi_m$ and (ii) the sets of phases related to each π -collection and referred to by the symbols F_1, \dots, F_n . The phases of activity are elements of the set of all neuronally possible phases N_i . The sets of phases are elements of N_p . The operations denoted by the symbols $\approx, \neq, \varepsilon$, and \wedge serve to build complex neuronal states from primitive entities. The set of all complex neuronal states constructible in \mathbf{N} is N_s . Superior “N” signifies that symbols or sequences of symbols in square brackets are interpreted in the algebra \mathbf{N} . Thus $\varphi_1^N, \varphi_2^N, \dots, \varphi_m^N$ are phases of activity; $F_1^N, F_2^N, \dots, F_n^N$ are sets that comprise the phases of related π -collections; and $\approx^N, \neq^N, \varepsilon^N, \wedge^N$ are operations. Instead of F_1, F_2, \dots , we will sometimes use more suggestive capital letters like the H, V, R, G, B, and Y of figure 2.

In \mathbf{N} there is only one fundamental operation: being *synchronous* with. It is referred to by the operation symbol \approx and relates phases of activity to each other:

$[\varphi_i \approx \varphi_j]^N$ is the state $[\varphi_i$ is *synchronous with* $\varphi_j]^N$.

The remaining \mathbf{N} -operations are derivationally defined by means of standard symbols, with “ \neg ” and “ $\&$ ” signifying negation and conjunction, “ \exists ” the existential quantifier, “ x ” a variable, “(“ and “)” prentices, “ \in ” set membership. We can thus define *asynchrony* \neq^N in a natural way:

$[\varphi_i \neq \varphi_j]^N$ is the state $[\neg \varphi_i \approx \varphi_j]^N$.

If neurons of a π -collection, to which the set of phases F_j^N is assigned, show a certain phase of activity φ_i^N , we say that the phase φ_i^N or a synchronous equivalent is an element of the set F_j^N . To refer to this neuronal state, we define the relation of *pertaining* ε^N :

$[\varphi_i \varepsilon F_j]^N$ is the state $[(\exists x)(x \approx \varphi_i \ \& \ x \in F_j)]^N$.

A further operation is *co-occurrence* \wedge^N of two states p^N and q^N . It is trivially defined:

$[p \wedge q]^N$ is the state $[p \ \& \ q]^N$.

The four operations are motivated by the hypothesis of indicativity and synchrony. They allow us to give an algebraic description of the scheme shown in figure 2. Assuming that the middle-shaded neurons show the phase of activity φ_1^N , the dark-shaded neurons the phase φ_2^N and the light-shaded neurons the phase φ_3^N , figure 2 expresses the cortical state:

$[\varphi_1 \varepsilon V \wedge \varphi_1 \varepsilon B \wedge \varphi_2 \varepsilon V \wedge \varphi_2 \varepsilon R \wedge \varphi_3 \varepsilon H \wedge \varphi_3 \varepsilon B]^N$.

Language

The notation already suggests that the algebra \mathbf{N} might be isomorphic to a compositional and systematic language \mathbf{L} . Since languages can be treated as algebras, we may define:

$\mathbf{L} = \langle L_i, L_p, L_s; \varphi_1, \dots, \varphi_m; F_1, \dots, F_n; \approx, \neq, \varepsilon, \wedge \rangle$.

The entities of \mathbf{L} are indexical expressions like *this* and *that* (included in the set L_i), predicates like *red* and *vertical* (in L_p) and clauses like *this is red* or *this is the same as that* (in L_s). The primitive symbols $\varphi_1, \dots, \varphi_m$

pick out specific indexicals and the primitive symbols F_1, \dots, F_n specific predicates. Again we will sometimes use more suggestive capital letters instead of F_1, \dots, F_n . The fundamental operation of \mathbf{L} is *sameness* \approx^L :

$[\varphi_i \approx \varphi_j]^L$ is the clause $[\varphi_i$ is *the same as* $\varphi_j]^L$.

The remaining operations can derivationally be defined. *Difference* \neq^L :

$[\varphi_i \neq \varphi_j]^L$ is the clause $[\neg \varphi_i \approx \varphi_j]^L$.

Using ε as the symbol for predication this time, the *copula* ε^L , which links an indexical expression φ_i^L to a predicate F_j^L , is defined by:

$[\varphi_i \varepsilon F_j]^L$ is the clause $[(\exists x)(x \approx \varphi_i \ \& \ x \in F_j)]^L$.

The copula (English: “is”) enables us to paraphrase natural language sentences like *this is vertical* in \mathbf{L} : $[\varphi_1 \varepsilon V]^L$. The *conjunction* \wedge^L between two clauses p^L and q^L is defined:

$[p \wedge q]^L$ is the clause $[p \ \& \ q]^L$.

The sentence *there is a blue vertical, a red vertical, and a blue horizontal object* can now be paraphrased:

$[\varphi_1 \varepsilon V \wedge \varphi_1 \varepsilon B \wedge \varphi_2 \varepsilon V \wedge \varphi_2 \varepsilon R \wedge \varphi_3 \varepsilon H \wedge \varphi_3 \varepsilon B]^L$.

Isomorphism and Preserved Constituency

To prove that the algebras \mathbf{N} and \mathbf{L} are isomorphic, a number of conditions have to be warranted. (i) There are as many phases of activity in \mathbf{N} as there are indexical terms in \mathbf{L} . (ii) Each π -collection, respectively, each related set of phases in \mathbf{N} is assigned to exactly one predicate of \mathbf{L} . (iii) \mathbf{L} -clauses, by stipulation, are identical if and only if they are logically equivalent. For, cortical states are identical if and only if they are referred to by logically equivalent \mathbf{N} -descriptions. To ensure this non-trivial condition, we thus have to accept that order is irrelevant in \mathbf{L} . This leads to a non-standard notion of language: Concatenation, no longer, is the fundamental operation of concept composition. (iv) The two fundamental operations synchrony and sameness are isomorphic. If so, this isomorphism then conveys to all operations that have recursively been defined. Since sameness is a reflexive, symmetric, and transitive relation, we have to define synchrony between phases as a reflexive, symmetric, and transitive relation, too. This is consistent with recent neurobiological data (cf. Eckhorn, 2000) and the computer simulations of oscillatory networks mentioned above.

In previous sections we argued that an architecture might not be compositional even if it is syntactically homomorphic to a compositional language. To preserve semantic compositionality, the isomorphism between \mathbf{L} and \mathbf{N} must, in addition, preserve the constituent structure of the language. If a primitive term is a constituent of a complex term, the isomorphic counterpart of the primitive term must be a constituent of the isomorphic counterpart of the complex term. The primitives of \mathbf{L} are the indexicals $\varphi_1^L, \dots, \varphi_m^L$ and the predicates F_1^L, \dots, F_n^L . Every \mathbf{L} -operation will lead to targets with those primitives as constituents. The clause $[\varphi_1 \approx$




Figure 3: Predicted neuronal representation of relations. The state $[\varphi_1 \in G \wedge \varphi_2 \in R \wedge \langle \varphi_1, \varphi_2 \rangle \in \text{IN}]^N$ is shown. The phases φ_1 of the G-neurons (light shading) occurs on the IN-module only in superposition with the phase φ_2 of the R-neurons (dark-shading) forming the duplex phase $\{\varphi_1, \varphi_2\}$ (hybrid shading). Since φ_2 also occurs as simplex on the IN-module, the situation on the IN-module is rendered by $[\{\{\varphi_1, \varphi_2\}, \{\varphi_2\}\} \in \text{IN}]^N$. By definition, this is equivalent to $[\langle \varphi_1, \varphi_2 \rangle \in \text{IN}]^N$.

φ_2^L can impossibly be tokened without tokening the indexicals φ_1^L and φ_2^L . With respect to constituency, what is true for L is also true for N: The state $[\varphi_1 \approx \varphi_2]^N$ is tokened just in case the phases φ_1^N and φ_2^N are tokened. Two phases are synchronous only if both of them actually occur within the cortex. The same is true *mutatis mutandis* for asynchrony. In L, the indexical φ^L and the predicate F^L are constituents of the clause $[\varphi \in F]^L$. Therefore, the phase φ^N and the π -collection to which the set F^N relates must be tokened, whenever the cortex is in $[\varphi \in F]^N$. This is obviously true because φ^N cannot pertain to the π -collection unless both the phase and the π -collection occur in the cortex. Figure 2 illustrates that the isomorphism preserves constituent relations for every operation: The complex state shown can only be tokened if, indeed, certain bursts of activity and certain collections of neurons are tokened. We may infer that oscillatory networks are not only syntactically, but also semantically compositional.

Relations

The representation of relations poses a binding problem of second order. The sentence *this red vertical object is in that green horizontal object* not only binds four property representations into two object representations, it moreover binds the two object representations by the relation *in*. The constituency preserving isomorphism between L and N straightforwardly generates a prediction of how to realize relational representation by oscillatory networks: After L has been extended by the tools for representing relations known from logic, N has to be extended in a way that perpetuates the isomorphism and the congruence with respect to constituency structure. The tools needed in the extensions of L and N are the operation of pairing, a higher-order copula and relation constants, or, respectively, their neuronal counterparts. Following Kuratowski (1967), or-

dered *pairs* are by common standards defined as asymmetric sets of second order:

$$[\langle \varphi_i, \varphi_j \rangle]^{L/N} \stackrel{\text{def}}{=} [\{\{\varphi_i, \varphi_j\}, \{\varphi_j\}\}]^{L/N}.$$

With the relations R_1^L, \dots, R_k^L being sets of pairs, the higher-order copula links pairs to relations in the manner of set membership. On the neuronal level, the R_1^N, \dots, R_k^N can be interpreted as relational modules:

$$[\langle \varphi_i, \varphi_j \rangle \in R_i]^{L/N} \stackrel{\text{def}}{=} [\langle \varphi_i, \varphi_j \rangle \in R_i]^{L/N}.$$

The sentence *this green object is in that red object* can now be paraphrased in the extension of L:

$$[\varphi_1 \in G \wedge \varphi_2 \in R \wedge \langle \varphi_1, \varphi_2 \rangle \in \text{IN}]^L.$$

Its neuronal counterpart – superior “L” is replaced by superior “N” – is shown in figure 3. To achieve a distribution of phases thus complex, some neurons are required to show a superposition of two phases. The presented model, therefore, predicts multiplex activity as a means of representing relations. Gasser and Colunga’s (1998) simulation, which also uses superposed phases in relational representations, supports the prediction.

Neither Connectionism nor Classicism

Cognitive architectures can be distinguished along three features:

Syntactic Trees. There are mappings from ordered sets of argument representations onto target representations.

Constituency (presupposes trees). For every syntactic tree, its argument representations are constituents of its target representation.

Order (presupposes constituency). For every target representation, there is a determinate order among its constituents.

These features are each realized by every standard language: There is a syntax, words are constituents of phrases, and the words follow a determinate word order. We can now ask which of these features a certain cognitive model implements. Turing-style computers typically implement all three features because they build complex representations from primitive representation by concatenation following certain syntactic rules. Integrated connectionist/symbolic architectures only implement syntactic trees. They do not implement the principle of constituency and the principle of order. Oscillatory networks, however, implement both syntactic trees and the principle of constituency. They do not implement an order among representations.

Oscillatory networks lie in some sense in between classical and connectionist architectures. They resemble connectionist networks in many respects: They may serve as associative, content addressable memories. They process information in parallel. They are able to learn from examples. They degrade gracefully. Etc. Still, oscillatory networks are stronger than traditional connectionist networks because, in oscillatory networks, primitive representations are constituents of complex representations. The primitive representations

inherit their causal properties to complex representations and, thereby, determine their semantic properties. Oscillatory networks unite the virtues and avoid the vices of classical and connectionist networks. They are semantically compositional and systematic.

Acknowledgments

Research for this paper was sponsored by the National German Scholarship Foundation. It was enabled by a one-year research scholarship at Rutgers University and the Rutgers Center of Cognitive Science. I owe many of the presented insights to seminars and discussions with Andreas Engel, Thomas Metzinger, Wolf Singer, Jerry Fodor, Ernie LePore, Brian McLaughlin, Bruce Tesar and Gerhard Schurz. I am grateful, also, to the Berlin Colloquium "Philosophy Meets Cognitive Science".

References

- Barnden, J. A. (1991). Encoding complex symbolic data structures with some unusual connectionist techniques. In J. A. Barnden & Pollack, J. B. (eds.). *Advances in Connectionist and Neural Computation Theory Vol. 1: High-level connectionist models*. Norwood, NJ: Ablex Publishing Corp.
- Eckhorn, R. (2000). Neural mechanism of scene segmentation (abstract). *Symposium on Neural binding of space and time*. Leipzig: unpublished.
- Engel, A. K., König, P., & Singer, W. (1991). Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences, USA*, 88, 9136-40.
- Fodor, J. A. & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.
- Fries, P., Roelfsema, P. R., Engel, A. K., König, P., & Singer, W. (1997). Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proceedings of the National Academy of Sciences, US*, 94, 12699-12704.
- Gasser, M. & Colunga, E. (1998). *Where Do Relations Come From?* (Tech. Rep. 221). Bloomington, IN: Indiana University Cognitive Science Program.
- Gawne, T. J., Kjaer, T. W., & Richmond, B. J. (1996). Latency: Another potential code for feature binding in striate cortex. *Journal of Neurophysiology*, 76, 1356-1360.
- Gigerenzer, G. & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition*, 43, 127-171.
- Gray, C. M. & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences, USA*, 86, 1698-702.
- Gray, C. M., König, P., Engel, A. K. & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334-37.
- Horgan, T. & Tienson, J. (1996). *Connectionism and the Philosophy of Psychology*. Cambridge, MA: The MIT Press.
- Hubel, D. H. & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160, 106-154.
- Hubel, D. H. & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215-243.
- Hummel, J. E. & Holyoak, K. J. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427 - 466.
- Kreiter, A. K. & Singer W. (1992). Oscillatory neuronal responses in the visual cortex of the awake macaque monkey. *European Journal of Neuroscience*, 4, 369-75.
- Kuratowski, K. (1967). *Set Theory*. Amsterdam: North-Holland.
- Lange, T. E. & Dyer, M. G. (1989). High-level inferring in a connectionist network. *Connection Science*, 1, 181 -217.
- Lowel, S., & Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255, 209-12.
- O'Keefe, J. & Recce, M. (1993). Phase relationship between hippocampal place units and the hippocampal theta rhythm. *Hippocampus*, 3, 317-330.
- Rockland, K. & Lund, J. S. (1983). Intrinsic laminar lattice connections in primate visual cortex. *Journal of Comparative Neurology*, 216, 303-18.
- Schillen, T. B. & König, P. (1994). Binding by temporal structure in multiple feature domains of an oscillatory neuronal network. *Biological Cybernetics*, 70, 397-405.
- Schwartz, C. & Bolz, J. (1991). Functional specificity of the long-range horizontal connections in cat visual cortex: a cross-correlation study. *Journal of Neuroscience*, 11, 2995-3007.
- Shastri, L. & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, 16, 417-94.
- Singer, W. & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18, 555-86.
- Smolensky, P. (1995). Connectionism, constituency and the language of thought. In Macdonald, C., & Macdonald, G. (Eds.), *Connectionism*. Cambridge, MA: Blackwell.
- Sougné, J. P. (1999). *INFERNET: A Neurocomputational Model of Binding and Inference*. Doctoral Thesis, Université de Liège.
- Wehr, M. & Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neuronal assemblies. *Nature*, 284, 163-66.