

Implementing the (De-)Composition of Concepts: Oscillatory Networks, Coherency Chains and Hierarchical Binding

Markus Werning* and Alexander Maye†

* Department of Philosophy, Heinrich-Heine University, Universitätsstr. 1, 44025 Düsseldorf, Germany
werning@phil.uni-duesseldorf.de (corresponding author)

† Zuse Institute Berlin, Takustr. 7, 14195 Berlin, Germany
maye@zib.de

Abstract

The paper introduces oscillatory networks as a model of the realization of lexical and non-lexical complex concepts in the cortex. The network has both perceptual and semantic capabilities. Three adequacy conditions, the compositionality of meaning, the compositionality of content, and the co-variation with content, are satisfied. Coherency chains and hierarchical mechanisms of binding are discussed.

Introduction

The neuronal structure of the cortex is usually perceived as radically different from the semantic structure on which cognition is defined. This paper argues that, in spite of these *prima facie* differences, semantic structure can be reduced to neural structure if one assumes that oscillatory networks be an appropriate model of neural reality. Several neurobiological data on neural synchronization support this assumption (Singer & Gray 1995). In the context of this paper, semantic structures are assumed to subservise two purposes. First, its elements evaluate linguistic expressions semantically and can thus be regarded as the *meanings* of those expressions. Second, its elements are themselves semantically evaluable with respect to external *content*. The linguistic expression ‘dog’, e.g., has as meaning the concept [dog], which is some internal state of the cognitive system whatever it may reduce to neurally. This internal state is itself a representation and must thus have external content. Its external content will here be identified with the set of dogs and is thus identical to the denotation of the predicate ‘dog’. Since representations are supposed to co-vary reliably with what they represent (Fodor 1992), the latter identification requires that internal states reliable *co-vary* with the denotations of those expressions whose meanings they are.

Taking co-variation as a condition for internal representation seriously, however, means that any system with semantic states must have some sort of perceptual, or more modestly speaking, informational access to its environment. Otherwise an adequate co-variation relation between representation and content cannot be established. Unlike other

attempts of mapping semantic structures onto connectionist networks (e.g., Smolensky 1995; Shastri & Ajjanagadde 1993; for a review see Werning 2001) in a compositional way, the approach pursued in this paper fulfills the condition of co-variation because the network proposed here does have (weakly) perceptual capabilities.

In this paper we will consider two ways of achieving conceptual complexity. While it is obvious that (i) a non-lexical concept like [red vertical] is complex — it is composed from the concepts [red] and [vertical] —, it is less obvious (and more controversial) if (ii) some lexical concepts like [elephant] are complex, i.e., if they can be decomposed into less complex concepts. A potential decomposition of the concept [elephant] might, e.g., involve concepts of prototypical attributes of elephants, e.g., [animate], [big], [gray], or concepts of prototypical parts of elephants, such as [trunk], [tusks], [ears], [tail]. In the first part of the paper we will focus on the neural implementation of the composition of primitive attributive concept into complex non-lexical concepts, whereas in the second part we will introduce a neural model of the decomposition of complex lexical substance concepts.

A main working hypothesis of ours is that some lexical concepts, i.e., concepts expressible by syntactically unstructured expressions (words), are indeed complex, most importantly those that are expressed by concrete nouns and denote substances (i.e., individuals like mama, kinds like mice, or stuffs like milk, which may change their attributes, but can still be re-identified over time, cf. Millikan 1998). Concepts for attributes like [red] or [vertical], in contrast, may well be primitive, i.e., not further decomposable.

Another working hypothesis is that the distinction between primitivity and complexity on the conceptual level corresponds to a distinction between locality and distributivity on the neural level. We thus expect the neural correlates of primitive (attributive) concepts (e.g., for color, orientation) to be relatively local. In fact, cortical areas for about 30 attribute dimensions could be anatomically located in the monkey (Felleman & van Essen 1991). In contrast, the correlates of complex (substance) concepts, we assume, are highly distributed neural states. Complex concepts like [mama], [mouse] or [milk] are thus not expected to be realized by single cells — so-called grandmother cells —, but by cell assemblies that may pertain to highly distinct parts

of the cortex.

The third working hypothesis, finally, assumes that the mechanism that binds together the distributed neural activities in the case of complex concepts is the intra-cortical synchronization of the electrical discharges of neurons (Engel *et al.* 1990). So far this mechanism has been employed mainly to explain how the representations of attributes of one and the same object can be bound together to form the representation of the object with its attributes. What we envisage, however, is that the mechanism of neural synchronization is general enough to also bind together the representations of the parts of an object to form the representation of the object with its parts. Since the parts of an object themselves are represented by neuronal assemblies, the mechanism that binds together parts into a whole has to be of higher order: It has to bind various assemblies, each of which consists of synchronized neurons, into larger ‘meta-assemblies’. The second-order binding mechanism, however, has to be such that the temporally induced unity of each partial assembly will be preserved.

Compositionality

It is widely assumed that speaking of semantic evaluation implies that the evaluated structure, which is usually called a syntax, is *compositional* with respect to the evaluating structure, which is usually called a semantics (Janssen 1997). Compositionality now means that the semantic value of an element of the syntax is a structure-dependent function of the semantic values of its syntactic parts. Formally, this can be put as follows (Hodges 2001):

Definition 1 (Compositionality). *Let the pair $\langle T, \Sigma \rangle$ be a syntax, i.e., a structure that comprises a carrier set T of syntactic elements (expressions, concepts, etc.) and a set Σ of syntactic operations (conjunction, predication etc.), which are defined as partial functions from Cartesian products of T into T . Furthermore, let μ be a function of semantic evaluation from T into a set of semantic values (meanings, contents, denotations, etc.). Then μ is called compositional just in case, for any n -ary syntactic operation $\sigma \in \Sigma$, there is a function μ_σ such that, for every $\sigma(t_1, \dots, t_n) \in T$, the following equation holds:*

$$\mu(\sigma(t_1, \dots, t_n)) = \mu_\sigma(\mu(t_1), \dots, \mu(t_n)). \quad (1)$$

Assume we are given a first order predicate language with identity $PL^= = \langle L, \Sigma_L \rangle$. The set of syntactic operations Σ_L shall comprise the standardly defined syntactic operations of that language and the denotations of expressions in L be given by a compositional function $\nu : L \rightarrow W$, where W is some worldly structure of objects, sets of objects, etc. To be a representational system, our network should then be such that it realizes a neuronal structure $\mathcal{N} = \langle N, \Sigma_N \rangle$ that satisfies the following three adequacy conditions:

1. \mathcal{N} is a compositional semantics of meanings for $PL^=$, i.e., there is a compositional function of semantic evaluation $\mu : L \rightarrow N$.
2. \mathcal{N} is itself compositionally evaluable with respect to content, i.e., there is a compositional function of semantic evaluation $\kappa : N \rightarrow W$.

3. The elements of \mathcal{N} reliably co-vary with their contents, the elements of W , such that, for every expression $t \in L$, $\nu(t) = \kappa(\mu(t))$.

Oscillatory Networks

Von der Malsburg’s (1981) supposition that the synchronous oscillation of neuronal responses constitutes a mechanism that binds the responses of feature specific neurons when these features are instantiated by the same object has been frequently applied to explain the integration of distributed responses. Object-related neuronal synchrony has been observed in numerous cell recording experiments (reviewed by Singer 1999) and experiments related to attention (Steinmetz *et al.* 2000), perception (Fries *et al.* 1997) and expectation (Riehle, Grün, & Aertsen 1997). Those data suggest the hypothesis that the neural basis of object concepts are oscillation functions and that the neural basis of predicate concepts are clusters of feature-specific neurons.

From Gestalt psychology the principles governing object concepts are well known. According to two of the Gestalt principles, spatially proximal elements with similar features (similar color/ similar orientation) are likely to be perceived as one object or, in other word, represented by one and the same object concept. This principle is implemented in oscillatory networks by the following mechanism: Oscillators that select input from proximal stimulus elements with like properties tend to synchronize while oscillators that select input from proximal stimulus elements with unlike properties tend to de-synchronize. As a consequence, oscillators selective for proximal stimulus elements with like properties tend to form out synchronous oscillation functions when stimulated simultaneously. This oscillation can be regarded as one object concept. In contrast, inputs that contain proximal elements with unlike properties tend to cause anti-synchronous oscillations, i.e., different object concepts.

In our model a single oscillator consists of two mutually excitatory and inhibitory neurons, each of which represents a population of biological cells (Fig. 1). If the number of excitatory and inhibitory biological cells is large enough, the dynamics of each oscillator can be described by two variables x and y . They evolve over time according to the following differential equations:

$$\begin{aligned} \dot{x} &= -\tau_x x - g_y(y) + L_0^{xx} g_x(x) + I_x + N_x \\ \dot{y} &= -\tau_y y + g_x(x) - I_y + N_y. \end{aligned} \quad (2)$$

Here, τ_ξ ($\xi \in \{x, y\}$) are constants that can be chosen to match refractory times of biological cells, g_ξ are transfer functions, L_0^{xx} describes self-excitation of the excitatory cell population, I_ξ are external inputs and N_ξ white noise, which models fluctuation within the cell populations. With I_ξ above threshold, the solutions of (2) are limit-cycle oscillations (Maye 2003).

Oscillators are arranged on the three-dimensional grid forming a feature module. Two dimensions represent the spatial domain, while the feature is encoded by the third dimension. Spatially close oscillators that represent similar properties synchronize. The desynchronizing connections establish a phase lag between different groups of syn-




Figure 1: a) A single oscillator consists of an excitatory (x) and an inhibitory (y) neuron. Each neuron represents the average activity of a cluster of biological cells. I_0^{xx} describes the self-excitation of the excitatory neuron. I_x and I_y amounts to external input. b) Synchronizing connections (solid) are realized by mutually excitatory connections between the excitatory neurons and hold between oscillators within one layer. Desynchronizing connections (dotted) are realized by mutually inhibitory connections between the inhibitory neurons and hold between different layers. ‘R’ and ‘G’ denote the red and green channel. The scheme applies to further neighbors and the remaining channels as well. c) Oscillators are arranged in a 3D-topology. The shaded circles visualize the range of synchronizing (light gray) and desynchronizing (dark gray) connections of a neuron in the top layer (black pixel). There are no interactions between oscillators within a column.

chronously oscillating clusters. Feature modules for different feature dimensions, e.g., color and orientation, can be combined by establishing synchronizing connections between oscillators of different modules in case they code for the same stimulus region.

Stimulated oscillatory networks (e.g., by stimulus of Fig. 2a), characteristically, show object-specific patterns of synchronized and de-synchronized oscillators within and across feature dimensions. Oscillators that represent properties of the same object synchronize, while oscillators that represent properties of different objects de-synchronize. We observe that for each represented object a certain oscillation spreads through the network. The oscillation pertains only to oscillators that represent the properties of the object in question.

Semantics in the Hilbert Space Framework

An oscillation function $x(t)$ of an oscillator is the activity of its excitatory neuron as a function of time during a time window $[0, T]$. Mathematically speaking, activity functions are vectors in the Hilbert space $L_2[0, T]$ of in the interval $[0, T]$ square-integrable functions. This space has the inner product

$$\langle x(t)|x'(t) \rangle = \int_0^T x(t) x'(t) dt. \quad (3)$$

The degree of synchrony between two oscillations lies between -1 and $+1$ and is defined as

$$\Delta(x, x') = \frac{\langle x|x' \rangle}{\sqrt{\langle x|x' \rangle \langle x'|x' \rangle}}. \quad (4)$$

From synergetics it is well known that the dynamics of complex systems is often governed by a few dominating states, the eigenmodes. The corresponding eigenvalues designate how much of the variance is accounted for by that mode. The first four eigenmodes of the stimulated network are shown in Fig. 2b. The overall dynamics of the network is given by the Cartesian vector $\mathbf{x}(t)$ that contains the excitatory activities of all oscillators as components. The network state at any instant is considered as a superposition of the temporally constant, but spatially variant eigenmodes \mathbf{v}_i weighted by the corresponding spatially invariant, but temporally evolving characteristic functions $c_i(t)$ of Fig. 2c:

$$\mathbf{x}(t) = \sum c_i(t) \mathbf{v}_i. \quad (5)$$

The eigenmodes, for any stimulus, can be ordered along their eigenvalues so that each eigenmode can be signified by a natural number i beginning with the strongest.

The Hilbert space analysis allows us to interpret the dynamics of oscillatory networks in semantic terms. Since oscillation functions reliably co-vary with objects, they may be assigned to some of the individual terms a, b, \dots, x, y, \dots of the language $PL^=$ by the partial function α .

The sentence $a = b$ expresses a representational state of the system (i.e., the representation of the identity of the objects denoted by the individual terms a and b) to the degree the oscillation functions $\alpha(a)$ and $\alpha(b)$ of the system are synchronous. The degree to which a sentence ϕ of $PL^=$ expresses a representational state of the system, for any eigenmode i , can be measured by the value $d_i(\phi) \in [-1, +1]$. In

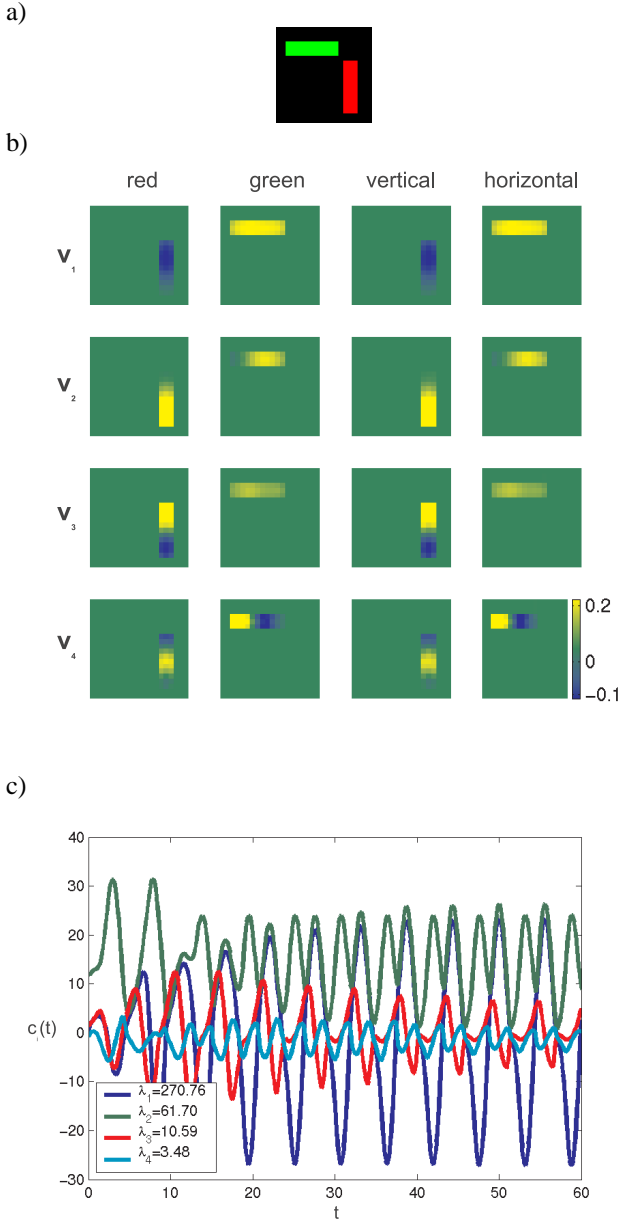


Figure 2: a) Stimulus: one vertical red bar and one horizontal green bar. b) First four eigenmodes. Each of the four eigenmodes $\mathbf{v}_1, \dots, \mathbf{v}_4$ with the largest eigenvalues is shown in one line. The four columns correspond to the four feature layers. Blue shading signifies negative, green zero and yellow positive components. c) The characteristic functions and eigenvalues for the first four eigenmodes.

case of identity sentences we have:

$$d_i(a = b) = \Delta(\alpha(a), \alpha(b)). \quad (6)$$

When we take a closer look at the first eigenmode of Fig. 2b, we see that most of the vector components are exactly zero (marked green). However, few components in the greenness and the horizontality layers are positive (marked yellow) and few components in the redness and the verticality layers are negative (marked blue). We may interpret this by saying that the first eigenmode represents two objects as distinct from one another. The representation of the first object is the positive characteristic function $+c_1(t)$ and the representation of the second object is the negative characteristic function $-c_1(t)$ (Because of the normalization of the Δ -function, only the signs of the eigenmode components matter). These considerations justify the following evaluation of non-identity:

$$d_i(\neg a = b) = \begin{cases} +1 & \text{if } d_i(a = b) = -1, \\ -1 & \text{if } d_i(a = b) > -1. \end{cases} \quad (7)$$

Feature layers function as representations of properties and thus can be expressed by predicates F_1, \dots, F_p , i.e., to every predicate F a diagonal matrix $\beta(F) \in \{0, 1\}^{k \times k}$ can be assigned such that, by multiplication with any eigenmode vector \mathbf{v}_i , the matrix renders the sub-vector of those components that belong to the feature layer expressed by F . To determine to which degree an oscillation function assigned to an individual constant a pertains to the feature layer assigned to a predicate F , we have to compute how synchronous it maximally is with one of the oscillations in the feature layer. We are, in other words, justified to evaluate the degree to which a predicative sentence Fa (read: ‘ a is F ’, e.g., ‘This object is red’) expresses a representational state of our system, with respect to the eigenmode i , in the following way:

$$d_i(Fa) = \max\{\Delta(\alpha(a), f_j) | f = c_i(t)\beta(F)\mathbf{v}_i\}. \quad (8)$$

If one, furthermore, evaluates the conjunction of two sentences $\phi \wedge \psi$ by the minimum of the value of each conjunct, we may regard the first eigenmode \mathbf{v}_1 of the network dynamics of Fig. 2 as a representation expressible by the sentence

This is a red vertical object and that is a green horizontal object.

We only have to assign the individual terms *this* ($= a$) and *that* ($= b$) to the oscillatory functions $-c_1(t)$ and $+c_1(t)$, respectively, and the predicates *red* ($= R$), *green* ($= G$), *vertical* ($= V$) and *horizontal* ($= H$) to the redness, greenness, verticality and horizontality layers as their neuronal meanings. Simple computation then reveals:

$$d_1(Ra \wedge Va \wedge Gb \wedge Hb \wedge \neg a = b) = 1. \quad (9)$$

Other stable eigenmodes (e.g., \mathbf{v}_2) represent alternative perceptual possibilities (e.g., the whole stimulus as one object).

Using further methods of formal semantics, the semantic evaluation of sentences has been extended to sentences comprising disjunction (\vee), implication (\rightarrow), and existential as well as universal quantifiers (\exists, \forall). The compositionality of meaning and the compositionality of content could

be proven (Werning 2004). Co-variation with content can always be achieved if the assignments α and β are chosen to match the network’s perceptual capabilities. Werning (2003b) extends this approach from an ontology of objects to an ontology of events. Werning (2003a) integrates relational concepts like [in]. Maye & Werning (2004) discuss ambiguous and illusionary representations. All (primitive or composed) concepts that can be expressed by a first order predicate language with identity are in principle compositionally realizable in our network, insofar as the lexical predicates of that language are semantically primitive. We will now turn to semantically complex lexical predicates.

The Decomposition of Complex Concepts: Coherency Chains and Hierarchical Binding

Among semanticists still some authors believe that no lexical concepts are decomposable (e.g., Fodor & Lepore 1992). According to those so-called atomist positions, only concepts that are linguistically expressible by syntactically combined expressions can be complex. In neuroscience, moreover, some authors still hold that substantial features like that of being an elephant, or even features as particular as that of being my grandmother are represented by highly specific single neurons (Logothetis & Sheinberg 1996).

The position we advocate is contrary to those, in two respects: We think that (i) some lexical concepts are decomposable into less complex concepts, and that (ii) the neural realization of those concepts is distributed over assemblies of neurons and meta-assemblies thereof. In psychology, philosophy and linguistics various theories have been proposed to account for the decomposition of concepts. The most prominent ones are: prototype theory (Rosch *et al.* 1976), lexical decomposition grammar (Hale & Keyser 1993; Wunderlich 1996), and frame theory (Minsky 1975; Barsalou 1999). For our purposes a modified version of frame theory seems to be most fruitful.

Frames are defined for large reference classes G of objects and allow for a categorization therein. Lowest-level frames can be rendered by a matrix with one row per attribute dimension (i) and the attributes F_{ij} of the respective dimension in the j -th column:

$$\text{dimensions} \begin{cases} \text{color} \\ \text{form} \\ \text{brightness} \\ \text{size} \end{cases} \underbrace{\begin{pmatrix} \text{red} & \text{green} \\ \text{round} & \text{square} \\ \text{light} & \text{dark} \\ \text{small} & \text{big} \end{pmatrix}}_{\text{frame}} \quad (10)$$

Relative to a frame, a category or concept C is rendered by a matrix that results from an assignment of typicality values $C_{ij} (\in [0, +1])$ to each attribute. C_{ij} tells how typical the j -th attribute, regarding the i -th dimension, is for an object that satisfies the concept C . Relative to the frame shown above, the matrix for the concept [cherry] might, e.g., look




Figure 3: Schematic view of a network that represents an object as a cherry relative to the frame 10. Each circle represents a layer of feature(F_{ij})-selective neurons. The lines mark the strongest chain of temporally coherent activity. Black/white dots represent high/low values for $C_{ij}d(F_{ij}x)$ (see equations 8 and 12).

as follows:

$$\begin{pmatrix} 1 & 0.2 \\ 0.9 & 0 \\ 0.9 & 0.9 \\ 1 & 0.1 \end{pmatrix} \quad (11)$$

If we assume that the list of attributes in each dimension are a partitioning of the reference class G , and if we choose an appropriate system of many-valued logic (Gottwald 2001), we can prove the following inequality for the fuzzy value $d(\in [-1, +1])$ of the existential claim $\exists x Cx$ (read: ‘There are C’s’, e.g., ‘There are cherries’):

$$d(\exists x Cx) \geq \max_{\pi \in \Pi} \min_{i=1}^n (C_{i\pi(i)}d(F_{i\pi(i)}x)). \quad (12)$$

$x \mapsto L_2[0, T]$

Here, C is the predicate expressing the concept C , Π is the set of all permutations of the indexes j , and n is the number of attribute dimensions in the reference frame. With regard to the evaluation of d , the mapping $x \mapsto L_2[0, T]$ ranges over all assignments of the term x to an oscillation in $L_2[0, T]$.

To the result of equation 12 we can directly apply the neuronal interpretation of the value $d(F_{ij}x)$, which has already been developed in equation 8. From these considerations we may infer the following hypothesis:

Hypothesis 1 (Coherency chains). *Provided that a concept is completely decomposable into attributive concepts according to its matrix C , the lowest boundary of the degree to which the network represents an object as an instance of the concept is given by the strength of the strongest (max) weighted coherency chain (weights: C_{ij}) that pertains to layers of neurons selective for the attributes (F_{ij}) of the frame in question. Here, any coherency chain is regarded just as strong as the weakest (min) weighted coherency among the activities of the feature-selective neurons involved (see Fig. 3).*

Some complex concepts, however, are not completely decomposable into attributive concepts, because they have characteristic parts. In that case hierarchical binding mechanisms are required (see Fig. 4). To explore their nature is in the focus of ongoing research.

Conclusion

Both, complex non-lexical concepts that can be composed from attribute concepts and complex lexical concepts that can be decomposed into attribute concepts can be modelled




Figure 4: Schematic view of the successive composition of the complex concept [traffic-light] relative to the frame 10. The first row shows the characteristic patterns of temporally coherent network activity for the concepts of the object parts. A higher-order binding mechanism stepwise composes the concept of the whole object from the concepts of its parts.

by oscillatory networks in a biologically realistic way. With regard to the latter, coherency chains are required. Further research needs to be done to account for complex lexical concepts that decompose into concepts of parts. Here, mechanisms of hierarchical binding are postulated.

References

- Barsalou, L. W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22:577–660.
- Engel, A. K.; König, P.; Kreiter, A. K.; and Singer, W. 1990. Inter-areal and inter-hemispheric synchronization of oscillatory responses in cat visual cortex. *Society of Neuroscience Abstracts* 16:1269.
- Felleman, D. J., and van Essen, D. C. 1991. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* 1:1–47.
- Fodor, J., and Lepore, E. 1992. *Holism: A Shopper's Guide*. Oxford: Blackwell.
- Fodor, J. 1992. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- Fries, P.; Roelfsema, P. R.; Engel, A. K.; König, P.; and Singer, W. 1997. Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc. Natl. Acad. Sci. USA* 94:12699–704.
- Gottwald, S. 2001. *A Treatise on Many-Valued Logics*. Baldock: Research Studies Press.
- Hale, K., and Keyser, S. 1993. On argument structure and the lexical expression of syntactic relations. In Hale, K., and Keyser, S., eds., *The View from Building 20*. Cambridge, MA: MIT Press. 53–109.
- Hodges, W. 2001. Formal features of compositionality. *Journal of Logic, Language and Information* 10:7–28.
- Janssen, T. 1997. Compositionality. In van Benthem, J., and ter Meulen, A., eds., *Handbook of Logic and Language*. Amsterdam: Elsevier. 417–73.
- Logothetis, N. K., and Sheinberg, D. L. 1996. Visual object recognition. *Annual Review of Neuroscience* 19:577–621.
- Maye, A., and Werning, M. 2004. Temporal binding of non-uniform objects. *Neurocomputing* 58–60:941–8.
- Maye, A. 2003. Correlated neuronal activity can represent multiple binding solutions. *Neurocomputing* 52–54:73–77.
- Millikan, R. G. 1998. A common structure for concepts of individuals, stuffs and real kinds: More Mama, more milk, and more mouse. *Behavioral and Brain Sciences* 21:55–100.
- Minsky, M. 1975. A framework for representing knowledge. In Winston, P. H., ed., *The Psychology of Computer Vision*. McGraw-Hill.
- Riehle, A.; Grün, S.; and Aertsen, A. 1997. Spike synchronization and rate modulation differentially involved in motor cortical functions. *Science* 278:1950–53.
- Rosch, E.; Mervis, C.; Gray, W.; Johnson, D.; and Boyes-Braem, P. 1976. Basic objects in natural categories. *Cognitive Psychology* 8:382–439.
- 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 Sciences* 16:417–94.
- Singer, W., and Gray, C. M. 1995. Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* 18:555–86.
- Singer, W. 1999. Neuronal synchrony: A versatile code for the definition of relations? *Neuron* 24:49–65.
- Smolensky, P. 1995. Reply: Constituent structure and explanation in an integrated connectionist/symbolic cognitive architecture. In Macdonald, C., and Macdonald, G., eds., *Connectionism*. Cambridge, MA: Blackwell. 223–290.
- Steinmetz, P. N.; Roy, A.; Fitzgerald, P. J.; Hsiao, S. S.; Johnson, K. O.; and Niebur, E. 2000. Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature* 404:187–90.
- von der Malsburg, C. 1981. The correlation theory of brain function. Internal Report 81-2, Max-Planck-Institute for Biophysical Chemistry, Göttingen.
- Werning, M. 2001. How to solve the problem of compositionality by oscillatory networks. In Moore, J. D., and Stenning, K., eds., *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*. London: Lawrence Erlbaum Associates. 1094–1099.
- Werning, M. 2003a. Synchrony and composition: Toward a cognitive architecture between classicism and connectionism. In Löwe, B.; Malzkorn, W.; and Raesch, T., eds., *Applications of Mathematical Logic in Philosophy and Linguistics*, Trends in Logic. Dordrecht: Kluwer. 261–278.
- Werning, M. 2003b. Ventral vs. dorsal pathway: the source of the semantic object/event and the syntactic noun/verb distinction. *Behavioral and Brain Sciences* 26:299–300.
- Werning, M. 2004. The temporal dimension of thought: Cortical foundations of predicative representation. *Synthese*. Forthcoming.
- Wunderlich, D. 1996. Models of lexical decomposition. In Weigand, E., and Hundsnurscher, F., eds., *Lexical Structures and Language Use*. Niemeyer. 169–83.