



## Synchrony in the Eye of the Beholder: An Analysis of the role of Neural Synchronization in Cognitive Processes

FRANK VAN DER VELDE\* and MARC DE KAMPS

*Cognitive Psychology Unit, University of Leiden, Wassenaarseweg 52, 2333 AK Leiden, The Netherlands (\*Corresponding author: Phone: (31) (0)715273637; Fax: (31) (0)715273783; E-mail: vdvelde@fsw.leidenuniv.nl)*

(Received: 13 November 2001; in final form: 29 May 2002)

**Abstract.** We discuss the role of synchrony of activation in higher-level cognitive processes. In particular, we analyze the question of whether synchrony of activation provides a mechanism for compositional representation in neural systems. We will argue that synchrony of activation does not provide a mechanism for compositional representation in neural systems. At face value, one can identify a level of compositional representation in the models that introduce synchrony of activation for this purpose. But behavior in these models is always produced by means conjunctive representations in the form of coincidence detectors. Therefore, models that rely on synchrony of activation lack the systematicity and productivity of true compositional systems. As a result, they cannot distinguish between type and token representations, which results in misrepresentations of spatial relations and propositions. Furthermore, higher-level cognitive processes will likely integrate information from widely distributed areas in the brain, which puts severe restrictions on the underlying neural dynamics if synchrony of activation is to play a role in these processes. We will briefly discuss these restrictions in the case of feature binding in visual cognition.

**Key words:** cognitive processes, compositional representations, conjunctive representations, neurodynamics, productivity, synchrony, systematicity

### 1. Introduction

A number of studies have reported the occurrence of synchronous neural activation in the brain during visual processing. In a characteristic experiment (e.g., Singer and Gray, 1995), one or more light bars are moved across the receptive fields of two groups of neurons. A synchrony of activation is found between both groups of neurons when a single light bar (i.e., a coherent object) is moved across the receptive fields of both groups of neurons. In contrast, synchrony of activation is not found when two light bars are moved in opposite directions, each one across the receptive field of one group of neurons. Synchrony of activation occurs when neurons produce activity (spikes) at the same moment in time (in phase). In particular, it is assumed that spike timing of the neurons should not differ more than 10

msec for synchrony of activation to occur (e.g., Gray, 1999; Singer, 1999; Von der Malsburg, 1999).

In the processing of visual information, synchronous activation of neurons could be used as a measure of the fact that the neurons are responding to the same object. In this manner, synchrony of activation could provide an answer to the so-called binding problem in visual perception, which refers to the question of how the activity of different neurons can be integrated as a representation of the same object (e.g., Singer and Gray, 1995). On the basis of this interpretation, synchrony of activation could play an important role in cognition in general, because binding problems occur in many (higher-level) cognitive processes.

In particular, binding problems occur in processes that rely on structural or compositional forms of representation (e.g., Hummel and Biederman, 1992; Bienenstock and Geman, 1995). A compositional representation consists of constituent representations as its parts and the relations between these parts. A compositional representation stands in contrast with a combinatorial or conjunctive representation, which is a dedicated representation for a specific conjunction of constituents. Examples of the latter are the so-called 'cardinal' cells (Barlow, 1972), which are neurons that respond selectively to specific conjunctions of visual features like color, form and texture.

A fundamental problem that occurs with conjunctive representations is the lack of systematicity and productivity in this form of representation (e.g., Fodor and Pylyshyn, 1988). For instance, one could have cardinal cells for the conjunctions *red-triangle*, *green-triangle* and *blue-square*, but not for the conjunctions *blue-triangle* and *red-square*. In this manner, one would be able to recognize red triangles and green triangles, without understanding that they are systematically related as triangles. Furthermore, one would have the ability to recognize red triangles and blue squares but not the ability to recognize blue triangles and red squares, which shows a lack of productivity in handling novel conjunctions of familiar colors and forms.

In contrast, compositional representations are systematic and productive. In the case of visual features, for instance, compositional representation entails that objects will be represented in terms of the constituent representations for features like colors and forms. A particular object, e.g., a red triangle, will then be represented by means of a temporary compositional representation of color with form. In this way, it is possible to recognize the systematic relation between red and green triangles, because the constituent representation for *triangle* is part of both compositional representations. Furthermore, it is possible to recognize any novel conjunction of the familiar colors and forms, because the novel conjunction is represented as yet another composition of the familiar constituent representations. For these reasons, compositional representations are in particular important for (higher-level) cognitive processes (e.g., Fodor and Pylyshyn, 1988).

However, a fundamental difficulty with compositional forms of representation is the correct (temporal) binding of the constituent representations and rela-

tions in a given compositional representation. The proposal at hand is that these binding problems can be solved in terms of synchrony of activation (Hummel and Biederman, 1992; Tononi, Sporns and Edelman, 1992; Shastri and Ajjanagadde, 1993, Hummel and Holyoak, 1997). Before evaluating this proposal, we will first describe a few examples in more detail.

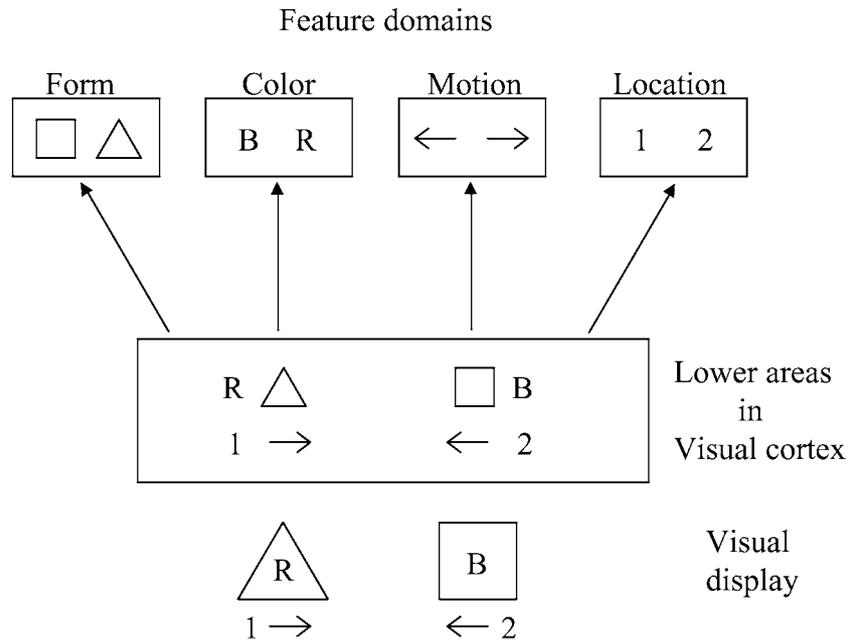
### 1.1. BINDING OF VISUAL FEATURES BY SYNCHRONY

It is now well established that different features (or attributes) of a visual object (e.g., its form, color, motion and location) are processed in different areas and pathways in the visual cortex (e.g., Farah, Humphreys and Rodman, 1999). When an object appears in the visual field, it will activate representations for each of its features in the different areas responsive to these features. With only one object in the visual field, each of these representations can be assumed to belong to the same object. But with more objects in the visual field, a cross-talk between representations can occur. Schematically, this is illustrated in Figure 1. A visual display of two objects, a red triangle and a blue square, is presented, with each object moving in a particular direction. Each object activates representations of its features in feature domains (or maps) for form, color, motion and location. On the level of the feature domains, however, it is now ambiguous whether the triangle or the square is red or blue, and on which location the square or the triangle, or the red object or the blue object, is placed, and in which direction it moves.

Neural synchronization could provide a mechanism for binding the features of an object into a coherent whole, if neurons representing the features of one object fire in synchrony with one another, but out of synchrony with neurons that represent features of other objects. In that case, the synchronous activation could be used to 'tag' or 'label' the features that belong to one object (e.g., Singer et al., 1997). Thus, if, in the feature domains in Figure 1, the representations for *red* and *triangle* fire in synchrony with one another, but out of synchrony with the representations for *blue* and *square*, this would represent the fact that one of the objects is a red triangle. Likewise, the synchronous activation of the representations for *blue* and *square* would indicate that the other object is a blue square. In the model of Tononi et al. (1992), form, color, motion and location are processed in parallel in separate pathways and areas, which are connected with re-entry connections. The model achieves synchronous activation (phase coherence) for features belonging to the same object.

### 1.2. BINDING OF SPATIAL RELATIONS BY SYNCHRONY

In the object recognition theory of Biederman (1987), a complex visual object is represented as the composition of particular elementary objects (geons) and their spatial relations. Thus, an object like a triangle on top of a square will be represented in terms of a compositional representation, which is constructed with



*Figure 1.* Activation of representations when a visual display contains two objects, a red triangle moving to the right and a blue square moving to the left. The objects first activate local feature representations in the lower areas of the visual cortex. In turn, these local representations activate feature representations in the feature domains for form, color, motion and location. In feature integration with synchrony of activation it is assumed that the feature representations of one object fire in synchrony with one another, and out of synchrony with the feature representations of other objects.

the representations for *triangle* and *square*, and the representations for the spatial relations *above* and *below*. The compositional representation will be correct if the representation for *triangle* is bound to the representation for *above*, and the representation for *square* is bound to the representation for *below*. Hummel and Biederman (1992) presented a model of visual object recognition in which an object is parsed into its constituent parts and their relations. Synchrony of activation is then used to temporarily bind the features and relations belonging to one object. Thus, in case of a triangle on top of a square, the unit for *triangle* fires in synchrony with the unit for *above*, and the unit for *square* fires in synchrony with the unit for *below*.

### 1.3. SYNCHRONY BEYOND VISION

Binding problems in cognition are not restricted to the binding of visual features and the binding of objects and spatial relations. In fact, the binding of objects and spatial relations is an example of binding arguments to predicates. In general,

arguments have to be bound to predicates to express propositions in terms of compositional representations. For instance, a compositional representation of *John loves Mary* is constructed out of the representations of *John*, *Mary* and the predicate *love(x,y)*. The slots  $(x,y)$  in the predicate *love(x,y)* represent the thematic roles of the verb *love*. That is, a verb like *love* demands a agent (subject) and a theme (object) in each proposition in which the verb occurs. Therefore, in the compositional representation of the proposition *John loves Mary*, the representations of *John* and *Mary* have to be bound to the thematic roles of agent and theme of the predicate *love(x,y)*. An essential feature of this form of binding is the relational content that has to be expressed by the binding mechanism. In the proposition *Mary loves John*, the representations of *John* and *Mary* are bound to the thematic roles of theme and agent respectively, which distinguishes this proposition from the proposition *John loves Mary*, even though the constituent representations *John*, *Mary* and *love(x,y)* are the same in both propositions.

Hummel and Holyoak (1997) presented a theory and computational model of analogical access and mapping. In their model, synchrony is used to bind arguments to predicate roles. For example, in the case of the proposition *John loves Mary*, there is a synchronous activation of the unit that represents *John* and the unit that represents the agent role of *love*, and a (different-phase) synchronous activation of the unit that represents *Mary* and the unit that represents the patient (theme) role of *love*. A similar form of argument to predicate binding is found in the model of reflexive reasoning presented by Shastri and Ajjanagadde (1993). Here, the fact *John gives Mary a book* is represented by the respective synchrony of activation between the units for the thematic roles of the predicate *give(x,y,z)* and the units for the objects *John*, *Mary*, and *book*.

#### 1.4. THE ROLE OF SYNCHRONY IN COGNITIVE PROCESSES

If synchrony of activation could provide a mechanism for solving binding problems in compositional forms of representation, synchrony would be an important neural mechanism for the generation of cognitive behavior. However, we will argue that synchrony of activation does not provide a neural mechanism for structural or compositional forms of representation.

As noted, a compositional representation is a complex representation that consists of primitive components (i.e., constituents like objects and relations) which are combined in a structural manner. At face value, one can identify a level of representation of this kind in the models that introduce synchrony of activation as a mechanism for compositional representation. However, we will show that, on its own, such a level of representation does not yet constitute a true form of compositional representation. For a representation to be truly compositional, it is necessary that the content of the representation can be analyzed in a compositional manner, so that, for instance, questions can be answered or behavior can be produced on the basis of the compositional representation.

For example, a compositional representation of a red triangle would consist of the constituent representations for *red* and *triangle*. Behavior in this case could consist of identifying the color when the form is known, or vice versa. In a true compositional representation of a red triangle, the color of this object will be identified on the basis of the constituent representation for *red*, and its form will be identified on the basis of the constituent representation for *triangle*. We will show that models based on synchrony of activation fail in this respect. They do not use the constituent representations to identify color or form, but instead use a conjunctive representation of *red-triangle* for this purpose. Likewise, in a compositional representation of *John loves Mary*, the question “Who does John love?” will be answered on the basis of the constituent representation for *Mary*, using the constituent representations for *John* and *loves* to analyze the sentence. In contrast, we will show that models based on synchrony of activation rely on a conjunctive representation of *John loves Mary* to answer questions of this kind.

The use of conjunctive representations in models based on synchrony of activation results from the fact that binding representations by means of the synchrony of their activation implies a mechanism that can detect the synchronous activation of these representations. In turn, this leads to the use of coincidence detectors that represent conjunctions of representations. Thus, contrary to the aim, the use of synchrony of activation in cognitive processes implies the use of conjunctive representations, which eliminates the possibility for true compositional forms of representation. In fact, synchrony of activation is not introduced to avoid conjunctive representations in these models, but to make effective use of them. That is, synchrony of activation is introduced to avoid the “superposition catastrophe” (von der Malsburg, 1987) that can occur with conjunctive forms of representation. In the following sections, we will discuss these issues in more detail.

## 2. The Need for Conjunctive Representations

One might suppose that binding by synchrony is directly achieved with the synchronous activation of the respective neural representations. Thus, in case of the binding problem illustrated in Figure 1, one might suppose that the feature representations of the red triangle in the feature domains are bound in an implicit manner because the respective neurons in these domains fire in synchrony. However, this is a naive (or ‘homunculus’) view of binding with synchrony, because it does not take into account the position from which the synchrony of activation is observed (e.g., see Dennett, 1991). From the perspective of an outside observer it is clear that the feature representations of the red triangle in Figure 1 carry the same label as a result of their synchronous activation, and thus belong to the same object. But from the perspective of the feature representations within the system this information is not directly available. The neurons that code for *red*, for instance, do not ‘know’ that they fire in synchrony with the neurons that code for *triangle*.

The fact that mere synchronization of activation is not enough to achieve binding is acknowledged by advocates of the binding-by-synchrony hypothesis. For instance, Mountcastle (1998, p. 376) clearly formulates the difficulties involved:

The binding hypothesis is set at a limited range, for it proposes no solution for the problem that bedevils many formulations of the neural mechanisms in perception: what next? What neural mechanisms can be imagined to be sensitive to and recognize the presence of synchronization versus its absence? What mechanism can identify – i.e., perceive – the pattern of synchronized system activity as that of a particular external event?

To answer this problem, Tononi et al. (1992) provided a behavioral response to evaluate their model. In their words (1992, p. 325):

One of the objectives of this work was to show that the patterns of activities and correlations that emerge in the model can be used for a discriminatory behavior. The presence of an output that can be independently evaluated eliminates potential ambiguities that might arise when one attempts to interpret these patterns in their own terms.

It is clear that the ‘potential ambiguities’ that can arise when activation patterns are interpreted ‘in their own terms’, refer to the naive or implicit view of binding by synchrony. An outside observer would see that the activation patterns achieved in the model are in synchrony for features of the same object. But the information available for an outside observer is not available within the system itself. However, if the model can generate an unambiguous behavioral response on the basis of the synchronous activation, then it is clear that information about synchrony is available within the system. The behavioral response used by Tononi et al. consisted of an eye movement to the location of a selected object (a red cross). The fact that the model could generate an eye movement to the location of the red cross indicates that it succeeded in binding information about the identity of that object with information about its location.

This procedure can be used as a test for binding. Thus, a system has achieved binding of representations if it can produce a response (discriminatory behavior) that depends on this binding. This discriminatory behavior can be seen as the answer to a particular ‘binding question’. In terms of the binding problem illustrated Figure 1, a model has achieved binding between form, color, motion and location if it can answer binding questions like “What is the color of the square?”, “What is the motion of the red object?”, “What is the location of the square?” or “What is the location and motion of the red triangle?”.

## 2.1. CONJUNCTIVE REPRESENTATIONS OF VISUAL FEATURES

Tononi et al. (1992) used operant conditioning in their model to produce an eye movement to the location of the selected object (red cross). Thus, whenever an eye

movement was made to the location of the red cross (by chance initially), a saliency system was activated that produced synaptic changes in the model on the basis of the activations produced by the red cross. Over time this produced strengthened connections between the feature representations for *red* (in the color area) and *cross* (in the form area) and the location representations in the location area and the motor output area. In other words, the model learned conjunctive representations of the features of the red cross and its potential locations. As a result of these conjunctive representations, the model could select the red cross from other objects and produce an eye movement to its location.

The need for coincidence detectors in this respect is clearly acknowledged in the literature. For instance, Singer et al. (1997) presented a review of synchronicity in which they refer to the existence of coincidence detectors as a necessary requirement for the synchrony hypothesis to work. In their words (p. 254, italics ours):

Two prerequisites need to be fulfilled in order to exploit synchronization as a selection and binding mechanism: firstly, *neurons must be able to act as coincidence detectors* and, secondly, mechanisms must exist that permit rapid and context-dependent synchronization.

Similar statements can be found in Gray (1999), Singer (1999), Shadlen and Movshon (1999) and Von der Malsburg (1999).

But coincidence detectors are conjunctive representations, which can be illustrated further with the binding problem presented in Figure 1. For instance, to answer the question “What is the color of the triangle?”, the synchrony between *red* in the color domain and *triangle* in the form domain has to be detected. The synchrony can be detected with a coincidence detector that is only activated when it receives synchronous activation from *red* and *triangle*. Historically, the fact that neurons could be coincidence detectors has been a major argument in favor of binding by synchrony (e.g., von der Malsburg, 1987). But a coincidence detector of *red* and *triangle* is also a conjunctive representation for *red triangle*.

Thus, synchrony of activation does not provide a mechanism for compositional representation of visual features, with which conjunctive feature representations could be avoided. In fact, the role of synchrony of activation is not to avoid conjunctive feature representations, but to use them effectively. For instance, without synchrony of activation the red triangle and the blue square in Figure 1 would activate four conjunctive representations for color and form, not only the correct conjunctions *red triangle* and *blue square*, but also the spurious conjunctions *red square* and *blue triangle*. With synchrony, this “superposition catastrophe” can be avoided because the conjunctive representations now consist of coincidence detectors, which precludes the spurious detection of a red square and a blue triangle. Historically, synchrony of activation was introduced for the purpose of solving the superposition catastrophe (Von der Malsburg, 1987). Over time, this issue became confused with the issue of compositional representation.

Summarizing, to answer questions about the binding of a set of features using the synchronous activation of their representations, coincidence detectors are needed that detect the synchrony of activation. The question to be answered determines the kind of coincidence detector, and thus the kind of conjunctive representation, that is needed. Questions like “What is the color of the triangle?” or “What is the location of the red object?” only require pairwise conjunctive representations of color with form and color with location respectively. But a question like “What is the location of the red triangle?” (e.g., TONI et al., 1992) requires a conjunctive representation of the features color, form and location, if synchrony of activation is used to answer the question.

## 2.2. CONJUNCTIVE REPRESENTATIONS OF SPATIAL RELATIONS

Hummel and Biederman (1992) presented a model in which synchrony of activation is used to generate a structural description of visual objects in terms of object parts (geons) and (spatial) relations between the object parts. For example, as illustrated in Figure 2, an object composed of a triangle above a square is represented by the model with units for the object parts *triangle* and *square* and the relative spatial relations *above* and *below*. Figure 2 gives a schematic presentation of the layers 3, 5, 6 and 7 of the model presented in Hummel and Biederman (1992). For instance, layer 3 in the model represents the whole set of attributes of the geons of an object. In Figure 2, this is simplified to a representation of the object parts *triangle* and *square*. Likewise, only the relative spatial relations *above* and *below* in layer 5 of the model are represented in Figure 2.

The primary theoretical objective of Hummel and Biederman was to show that the geon representations in layer 3 of the model can be synchronized with the appropriate relation representations in layer 5 of the model. In terms of Figure 2, this means that the representation for *triangle* will be in synchrony with the representation for *above*, and out of phase with the representation for *square*, which in turn is in synchrony with the representation for *below*.

However, as in the case illustrated in Figure 1, the claim that the synchrony of activation illustrated in Figure 2 represents binding between object parts and spatial relations is based on the naive or implicit view of binding. A true test of binding would be the ability to answer the relevant binding questions, i.e., to produce ‘discriminatory behavior’ (TONI et al., 1992). In this case, binding questions are “What object is above the square?” or “Where is the triangle relative to the square?”. Similar to the case illustrated in Figure 1, an answer to these questions would require the detection of the synchrony in activation between the representations for the object parts and the spatial relations. But this again requires conjunctive representations that operate as coincidence detectors.

The need for conjunctive representations in this case can be illustrated with the object classification performance of the model investigated by Hummel and Biederman (1992). To achieve object classification in this model, two additional

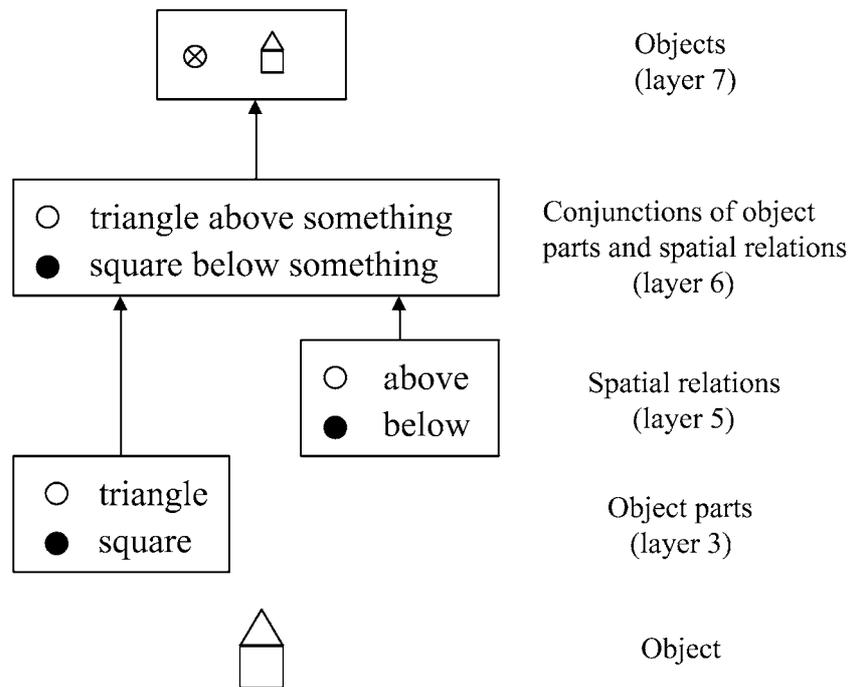


Figure 2. The use of synchrony of activation in the representation of spatial relations between object parts (geons) in an object, based on Hummel and Biederman (1992). An object activates units representing geons (layer 3) and spatial relations between geons (layer 5). The layer 3 and 5 units activate conjunctive representations (coincidence detectors) for the conjunctions of geons with spatial relations (layer 6). In turn, the layer 6 units activate the conjunctive representation for the whole object (layer 7). Units represented with filled circles fire in synchrony with one another and out of synchrony with units represented with open circles, which also fire in synchrony with one another. An 'x' unit fires in synchrony with both 'filled' and 'open' units.

layers (6 and 7) are introduced. In layer 6 conjunctions of object parts (geons) and relations are represented by means of coincidence detectors. Thus, in terms of Figure 2, there are units for *triangle above something* and *square below something*, with the unit for *triangle above something* in synchrony with the units for *triangle* and *above*, and the unit for *square below something* in synchrony with the units for *square* and *below*. In turn, layer 7 contains the units for whole objects, given by conjunctions of layer 6 units. Thus, in terms of Figure 2, layer 7 contains a unit for the whole object *triangle above square*. Because this unit represents the conjunction of *triangle above something* and *square below something*, it is in synchrony with both of these layer 6 units. In effect, this means that the units in layer 7 are not coincidence detectors, but simple conjunctive representations.

Hence, in the model of Hummel and Biederman (1992), synchrony does not generate a compositional representation of objects (object parts) and spatial rela-

tions, except on the basis of the naive or implicit view of binding by synchrony. Instead, as in the case of Figure 1, the importance of synchrony of activation in this model is not to avoid conjunctive representations but to use them effectively. That is, the purpose of synchrony of activation in this model is to avoid a superposition catastrophe for the units in layer 6. Without synchrony of activation, units in layer 6 would be activated for all possible conjunctions of object parts and spatial relations. With synchrony, only the units for the conjunctions *triangle above something* and *square below something* are activated in layer 6 of Figure 2.

As noted above, the units in layer 7 are not coincidence detectors, because they are activated by the units in layer 6, which fire out of synchrony. Therefore, a superposition catastrophe can occur with the units in level 7. This will occur when two or more objects are processed together. For instance, a second object in Figure 2, consisting of a diamond above a rectangle, would activate layer 6 units for *diamond above something* and *rectangle below something*. Together with the layer 6 units presented in Figure 2, this would not only result in the activation of layer 7 units for the true objects *triangle above square* and *diamond above rectangle*, but also for the spurious objects *triangle above rectangle* and *diamond above square*. For this reason, the model is restricted to process one object at a time, selected by spatial attention (Hummel and Biederman, 1992).

### 2.2.1. *Spatial Relations as Binary Predicates*

Although the need for conjunctive representations is similar for the cases illustrated in the Figures 1 and 2, there is a difference in the effectiveness of synchrony in both cases. This difference is related to the way in which predicates can be represented with synchrony of activation.

The binding of visual features, illustrated in Figure 1, can be described in terms of unary predicates. For instance, the binding of color and form in a red triangle can be described in terms of unary predicates such as ‘the triangle is red’, or ‘the red object is a triangle’. In general, a unary predicate  $P(x)$  expresses the fact that  $P$  belongs to or co-occurs with  $x$ . In turn, synchronous activation also expresses a form of co-occurrence. Hence, there is a similarity between the nature of unary predicates and synchronous activation, which motivates the use of coincidence detectors as conjunctive representations.

But in the case of higher order predicates, the similarity between the nature of the predicates and synchronous activation is lost. For instance, the binary relation *triangle above square* is different from the binary relation *square above triangle*, even though *triangle*, *square* and *above* co-occur in both instances. Yet, in terms of the synchronous activation of the representations *triangle*, *square*, and *above*, the distinction between *triangle above square* and *square above triangle* cannot be made. Hence, synchrony of activation is not suited to represent binary (and higher order) predicates in an explicit manner.

In the model of Hummel and Biederman (1992), binary predicates are represented implicitly, by means of single-place predicates such as *triangle above* and

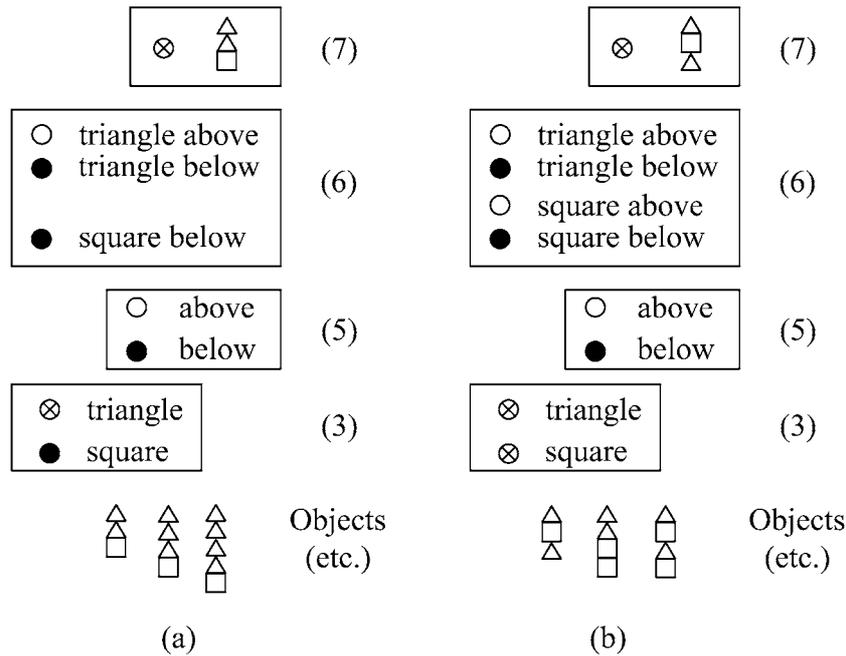


Figure 3. (a) Activations that occur in the model presented in Figure 2 with objects that consist of any number of triangles on top of a square. (b) Activations that occur in the model presented in Figure 2 with objects that consist of triangles intermixed with squares.

*square below*. The conjunctive representations (coincidence detectors) in layers 6 are representations of these single-place predicates, as illustrated in Figure 2. A whole object, such as a triangle above a square, is represented as a combination of single-place predicates, by means of the conjunctive representations in layer 7 (see Figure 2). However, when relations are not represented explicitly, difficulties in the expression of multiple relations will arise. Hummel and Biederman illustrate this fact with objects that consists of four geons on top of each other, in which case misrepresentations result for the relative positions of the two middle geons. But, in fact, serious misrepresentations can already arise in the case of objects that consist of two or three object parts on top of each other.

Figure 3a shows the representations that are activated when another triangle is placed on top of the object presented in Figure 2. In the layers 3 and 5, the object now activates the representation for *square* and *triangle* and the representations for *above* and *below*, with *square* in synchrony with *below* and *triangle* in synchrony with both *above* and *below*. In layer 6, this results in the activation of the conjunctive representations (coincidence detectors) for *square below*, *triangle above* and *triangle below*. Together, the layer 6 representations will activate the representation of the object in layer 7.

However, as illustrated in Figure 3a, the same configuration of activated representations will arise in the case of an object that consists of three (or more) triangles on top of a square. The reason is that in the case of three (or more) triangles on top of the square, again the representations for *square*, *triangle*, *above* and *below* are activated in layers 3 and 5, with *square* in synchrony with *below* and *triangle* in synchrony with both *above* and *below*. From that point onwards, the activations in the model are the same as in the case of two triangles on top of a square. In other words, with triangles on top of a square, all representations of object parts and spatial relations are already activated with two triangles. Adding more triangles to the object will not affect this pattern of activation, and thus will not be noticed by the model. In a similar manner, the model cannot distinguish between an object that consists of, say, two triangles on top of each other and objects that consist of three or more triangles on top of each other.

Figure 3b shows the pattern of activation that results when the square is in between two triangles. As before, the triangles produce the activation of *triangle* in synchrony with *above* and *below*. But *square* is now also in synchrony with *above* and *below*. This pattern of activation is not changed when more squares or triangles are added to the object. Thus, the model cannot distinguish between this object and any object that consists of two or more triangles combined with two or more squares, as illustrated in Figure 3b.

### 2.2.2. Types and Tokens in Structural Representations

The examples presented in the Figure 3 illustrate a fundamental difficulty that arises when synchrony of activation is used to generate structural or compositional representations. The representations (coincidence detectors) activated in the layers 3 and 5 are so-called type representations (e.g., Fodor and Pylyshyn, 1988). Each occurrence of a given object part, such as a triangle, activates the same *triangle* representation in layer 3. Likewise, each occurrence of a relative spatial relation, such as triangle above square, activates the same *above* representation in layer 5. Obviously, the conjunctive representations in layer 6 are also type representations, representing the conjunctions of object parts and relative spatial positions, because they are coincidence detectors of the synchrony of activation between the units in the layers 3 and 5. Thus, a second triangle on top of the object in Figure 2 results in the activation of the same unit for *triangle above something* in layer 6. Type representations are not sufficient to represent compositional structures. Instead, compositional forms of representation depend on so-called token representations, with which individual instances (or tokens) of the same type can be distinguished in a compositional structure (for examples, see van der Velde, 1997). The lack of token representation is evident in the examples presented in the Figure 3. It occurs for the representations of object parts, for the representations of spatial relations, and for the representations of the conjunctions of object parts and spatial relations. For instance, in case of the objects presented in Figure 3b, all type representations for *triangle above*, *triangle below*, *square above* and *square below* are activated.

This pattern of activation will not be affected by adding more tokens of these types to the object, as illustrated with the objects presented in Figure 3b.

### 2.3. CONJUNCTIVE REPRESENTATIONS BEYOND VISION

Hummel and Holyoak (1997) used synchrony of activation to bind arguments (objects) to predicate roles in propositional representations. Their model is in fact closely related to the object recognition model of Hummel and Biederman (1992).

Figure 4 illustrates the binding of objects to predicate roles of the proposition *John loves Mary* in terms of the model of Hummel and Holyoak (1997). As in Hummel and Biederman's (1992) model, there are units for objects and relations (predicate roles). The objects in this case are *John* and *Mary*. The predicate roles consist of the thematic roles of the predicate *love*, i.e., *love-agent* and *love-patient*. Synchrony of activation is used to represent role binding. Thus, in the case of *John loves Mary*, the unit for *John* fires in synchrony with the unit for *love-agent*, and out of synchrony with the unit for *love-patient*. In turn, the unit for *Mary* fires in synchrony with the unit for *love-patient*. The units for propositions, such as *John loves Mary* in Figure 4, fire in synchrony with all other units. Thus, similar to the units for whole objects in Figures 2 and 3, they are simple conjunctive representations. As a result, the unit for the proposition *John loves Mary* cannot be active together with units for other propositions, like *Susan loves Bill*. Otherwise, a superposition catastrophe on the level of proposition units would arise, resulting in confusions about who loves who. Therefore, only one proposition is active at the time in the model (Hummel and Holyoak, 1997), which is the equivalent of processing one object at the time in Figure 2.

As before, the mere fact that there is synchrony of activation between units does not entail that these units form a compositional representation (as assumed in the naive or implicit view of binding). A true form of compositional representation allows the generation of meaningful behavior without the use of conjunctive representations. In this case, that behavior would consist of answering binding questions like "What does John do?", "What happens to Mary?", or "Who does John love?" and "Who loves Mary?". However, as illustrated in Figure 4, in this case there are conjunctive representations for the conjunctions of objects with predicate roles, like *John love-agent* and *Mary love-patient*, which can be used to answer questions like "What does John do?" and "What happens to Mary?". But the answers to questions like "Who does John love?" or "who loves Mary?" in fact depend on a conjunctive representation for the whole proposition *John loves Mary*, which represents the conjunction of *John love-agent* and *Mary love-patient*. Hence, in this model as well, synchrony of activation is not used to avoid conjunctive presentations, but to use them effectively. Without synchrony, the model would activate the representations for all four conjunctions of objects and predicate arguments, which would represent four different proposition representations (i.e., *John/Mary loves John/Mary*).

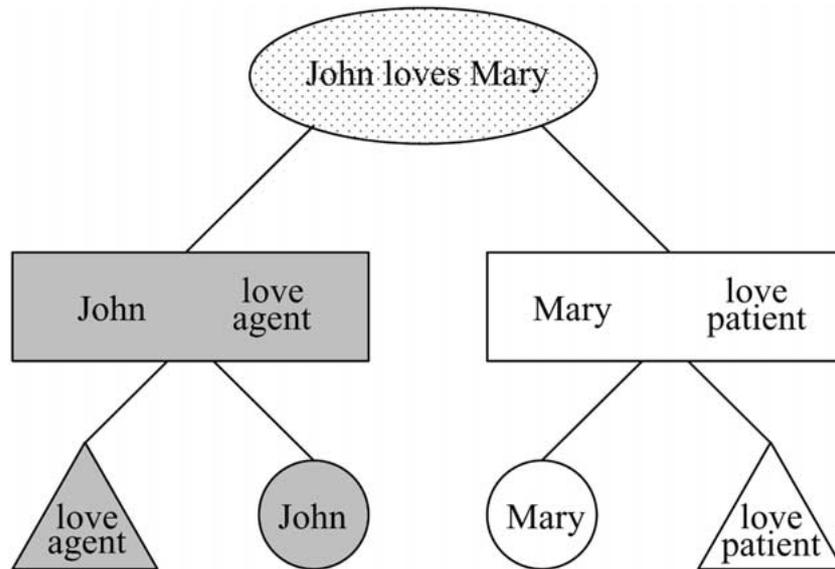


Figure 4. Representation of a proposition using synchrony of activation between units that represent objects (*John*, *Mary*) and predicate roles (*love-agent*, *love-patient*), based on Hummel and Holyoak (1997). The shaded units for the object *John*, the predicate role *love-agent*, and the conjunction *John love-agent* fire in synchrony with another, and out of synchrony with the synchronous white units for the object *Mary*, the predicate role *love-patient* and the conjunction *Mary love-patient*. The conjunctive representation for the proposition *John loves Mary* is in synchrony with the units for the conjunctions *John love-agent* and *Mary love-patient*.

The fact that synchrony of activation is not used to avoid conjunctive presentations, but to use them effectively is also clear in the model of reflexive reasoning presented by Shastri and Ajjanagadde (1993). In their model synchrony of activation is used to show how a known fact such as *John gives Mary a book* can result in an inference such as *Mary owns a book*. The proposition *John gives Mary a book* is represented by a coincidence detector (or ‘fact node’) that detects the respective synchrony in activation between the nodes for *John*, *Mary* and *book*, and the nodes for *giver*, *recipient* and *give-object*, that represent the thematic roles of the predicate *give(x,y,z)*. The reasoning process results in the respective synchronous activation of the nodes for *owner* and *own-object* of the predicate *own(y,z)* with the nodes for *recipient* and *give-object* of the predicate *give(x,y,z)*. As a result, the node for *Mary* is in synchrony with the node for *owner* and the node for *book* is in synchrony with the node for *own-object*. Hence, the proposition *Mary owns a book* can now be detected by a coincidence detector (fact node).

It is clear that the coincidence detectors (fact nodes) for the propositions *John gives Mary a book* and *Mary owns a book* are conjunctive representations of these propositions. Thus, these propositions are not represented in a compositional

manner. Instead, synchrony of activation is used to avoid the spurious activation of conjunctive representations for propositions like *Mary gives John a book* and *John owns a book*. These conjunctive representations are not activated because the node for *Mary* is not in synchrony with the node for *giver* and the node for *John* is not in synchrony with the nodes for *recipient* and *owner*.

However, as discussed with the examples presented in the Figure 3, using synchrony of activation to represent higher order predicates can easily result in misrepresentations. With synchrony of activation only unary or single-placed predicates can be represented explicitly in terms of conjunctive representations (by means of coincidence detectors). As a result, higher order predicates have to be represented implicitly as assemblies of these conjunctive representations. Because conjunctive representations are type representations, assemblies of conjunctive representations cannot represent multiple tokens of the same type.

Similar problems arise in the case of the model of reflexive reasoning of Shastri and Ajjanagadde (1993). For instance, a proposition like *John gives Mary a book and Mary gives John a pen* would cause confusion between *John* and *Mary* in their respective roles of *giver* and *recipient*. To solve this problem, Shastri and Ajjanagadde allowed for a duplication (or multiplication) of the representations for predicates. In this way, the whole proposition *John gives Mary a book and Mary gives John a pen* is partitioned into the two elementary propositions *John gives Mary a book* and *Mary gives John a pen*. Each elementary proposition is now represented with a specific node for the predicate  $give(x,y,z)$ . Thus, there are specific nodes for, say,  $give^1(x,y,z)$  and  $give^2(x,y,z)$ , with  $give^1(x,y,z)$  connected to the coincidence detector (fact node) for *John gives Mary a book* and  $give^2(x,y,z)$  connected to the coincidence detector (fact node) for *Mary gives John a pen*.

Notice that in this duplication solution, not only the representations for the predicates have to be duplicated, but also the associations between predicates that are to be used in the reasoning process. Thus, the node for  $give^1(x,y,z)$  has to be associated with a node for, say,  $own^1(y,z)$  and the node for  $give^2(x,y,z)$  has to be associated with a node for  $own^2(y,z)$ . A fundamental difficulty with this solution is the question of how these associations can be formed simultaneously during learning. During its development, a child will learn from specific examples. Thus, it will learn that, when *John gives Mary a book*, it follows that *Mary owns the book*. In this way, the child will form an association between the representations for  $give^1(x,y,z)$  and  $own^1(y,z)$ . However, in the duplication solution, the child would also have to form the association between the representations for  $give^2(x,y,z)$  and  $own^2(y,z)$ , even though these representations are not activated with *John gives Mary a book* and *Mary owns the book*.

This learning problem is closely related to the notions of systematicity and productivity. With the duplication of the representation for the predicate  $give(x,y,z)$  into  $give^1(x,y,z)$  and  $give^2(x,y,z)$ , the systematicity between *John gives Mary a book* and *Mary gives John a pen* is lost. As a result, the productivity in the use of the knowledge acquired with the association between  $give^1(x,y,z)$  and  $own^1(y,z)$  is lost

as well. That is, although the child may have learned that from *John gives Mary a book* it follows that *Mary owns the book*, it cannot use that knowledge productively in a novel situation, such as *Mary gives John a pen*, to deduce that *John owns the pen*. The lack of systematicity and productivity in this duplication solution is typical for conjunctive forms of representation.

Another problem with duplicating predicate representations is the nested nature of many propositions. For instance, a proposition like *John knows that Mary knows John* cannot be partitioned into two elementary propositions, because the elementary proposition *Mary knows John* is itself part of the predicate *know(John,y)*. Thus, as a result of synchrony of activation, confusions will arise between the representations of the predicates *know(John,y)* and *know(Mary,y)* in such a nested proposition.

#### 2.4. SUMMARY

We analyzed the question of whether synchronous activation of neurons or nodes can be used to implement compositional forms of representation. In each of the models we discussed, there is indeed a level (layer) of representation in which only constituents (objects and predicates) are represented, and in which conjunctions of constituents are represented in terms of synchrony of activation. However, these levels of representation do not constitute true forms of compositional representation, because (behavioral) decisions are not made on these levels. Instead, decisions, such as answers to binding questions, are always made by means of conjunctive representations, consisting of coincidence detectors that detect the synchrony of activation between the relevant constituent representations.

In other words, models based on synchrony of activation fail to produce true forms of compositional representation because they do not have the ability to de-compose a compositional representation into its constituent representations. For instance, in the case of a compositional representation for *John loves Mary*, the binding question “Who does John love?” is answered by analyzing or de-composing the compositional representation *John loves Mary* in such a manner that the constituent representation for *Mary* can be used to produce the answer to the question. Likewise, using a compositional representation to identify the color of a red triangle, the compositional representation *red-triangle* has to be decomposed so that the constituent representation for *red* can be used to identify the color of the object. Models based on synchrony of activation fail to decompose compositional representations, but instead use conjunctive representations to answer binding questions or to produce behavior.

Thus, synchrony of activation is not used to avoid conjunctive representations, as would be necessary in the case of compositional representation. Instead, synchrony is applied to make effective use of conjunctive representations. Without synchrony, spurious activations (a superposition catastrophe) of conjunctive representations would occur in the models we analyzed. Hence,

synchrony of activation is used to select the conjunctive representations that represent valid conjunctions of constituents.

However, synchrony of activation can only be used to select between conjunctive representations that represent unary or single-placed predicates, because synchrony of activation can only express co-occurrence, as found in unary or single-placed predicates. Binary or higher order predicates are only represented in an implicit manner with synchrony of activation. As a result, ambiguities can result even with synchrony of activation, as illustrated in Figure 3. Furthermore, even in the case of unary predicates, synchrony of activation will not necessarily result in the selection of the valid conjunctive representations, because the success of this selection will also depend on the underlying neural dynamics. We will discuss this further in the next section.

### 3. The Neural Dynamics of Synchrony

As noted in the introduction, for synchrony of activation to occur it is assumed that the spike timing of the neurons should not differ more than 10 msec (e.g., Gray, 1999; Singer, 1999; Von der Malsburg, 1999). This requires the existence of mechanisms that permit rapid and context-dependent synchronization (e.g., see the quote from Singer et al., 1999, section 2.1). Furthermore, the use of coincidence detection as a binding mechanism requires that the synchrony of activation is maintained throughout the whole process of activating the coincidence detectors. This requirement is in particular important in the case of higher-level cognitive processes, because these processes will often integrate information processed in widely distributed areas in the brain. We will briefly discuss this issue in the case of binding unary predicates, as illustrated in Figure 1.

The process of feature binding by means of synchrony of activation, illustrated in Figure 1, requires at least four restrictions on the underlying neural dynamics. First, the neurons that represent the local features of the red triangle in the lower areas of the visual cortex have to fire in synchrony with one another, but out of synchrony with the neurons that represent the local features of the blue square. Second, the neurons that represent *red* in the color domain have to fire in synchrony with one another, but out of synchrony with the neurons that represent *blue* in the color domain. A similar pattern of activation has to occur in each of the other feature domains. Third, the neurons that represent *red* in the color domain have to fire in synchrony with the neurons that represent *triangle* in the form domain. They also have to fire in synchrony with the neurons that represent *move-right* in the motion domain and the neurons that represent *location-1* in the location domain. Fourth, the activation from the unit representing *triangle* in the form domain has to arrive in synchrony with the activation from the unit representing *red* in the color domain at the site of the coincidence detector for *red triangle*, so that the conjunction between *red* and *triangle* can be detected. Moreover, the activation from, for instance, the unit representing *blue* should not arrive in synchrony with the activa-

tion from the unit representing *triangle* at the site of the coincidence detector for *blue triangle*, so that the spurious detection of a blue triangle is avoided.

Each of these restrictions is necessary to select the valid conjunction of features by means of synchrony of activation. For instance, if the activation delay between the form domain and the color-form coincidence detectors is different from the activation delay between the color domain and the color-form coincidence detectors, the synchronous activations of *triangle* and *red* in the form and color domains will arrive out of synchrony at the site of the color-form coincidence detectors, which would result in a failure to detect the conjunction *red triangle*. In fact, it could be that the activations of *triangle* and *blue*, which are out of synchrony at the level of the feature domains, arrive in synchrony at the site of the color-form coincidence detectors, which would result in the detection of the mistaken conjunction of *blue triangle*.

In the visual cortex, information about the location and the motion of an object is processed in the dorsal pathway, which leads to the parietal cortex. In contrast, form and color information are processed in the ventral pathway, which leads to the temporal cortex. Differences between activation delays in these pathways can easily occur even though both pathways originate from the same area (V1). Thus, even if there is synchrony of activation within each of the feature domains, these activations could be out of synchrony between the feature domains. There is indeed substantial evidence for large differences in the time course of activation in the visual cortex (e.g., see Ghose and Maunsell, 1999). In particular, the onset latencies (initial activations) are much shorter in the dorsal pathway compared to the ventral pathway (e.g., Schroeder, Mehta and Givre, 1998; Schmolesky et al., 1998; Mehta, Ulbert and Schroeder, 2000). For instance, Schmolesky et al. (1998) found that the average initial activation in the areas MT and MST of the dorsal pathway occurred about 6 to 9 msec after the average initial activation in the primary visual cortex (V1). In contrast, the average initial activation in area V4 of the ventral stream, which is halfway V1 and the final stage of the ventral pathway (IT), lagged 20 to 40 msec (or more) behind the initial activation in V1. Furthermore, the distribution in onset latency was much larger in the ventral pathway compared to the dorsal pathway. Schroeder, Mehta and Givre (1998) and Mehta, Ulbert and Schroeder (2000) found similar differences between these and other areas in the dorsal and ventral pathway.

As these results show, there are significant differences in onset latencies between the dorsal and ventral pathway, and the distribution in onset latency is also much larger in the ventral pathway compared to the dorsal pathway. The differences in activation delay reported in the literature are well beyond the limit of 10 msec, needed for synchrony of activation to occur (Gray, 1999; Singer, 1999; Von der Malsburg, 1999). To illustrate this in terms of Figure 1, binding by synchrony of activation would require that the feature representations for the red triangle in the feature domains fire in synchrony with one another within a time window of about 10 msec (and out of synchrony with the feature representations for the blue

square). To achieve this, the difference in onset latency of activation between, say, the feature representations of form and motion would have to be compensated for. This could occur if this difference in onset latency would match the oscillation phase by which the neurons fire in these domains. However, given the large variability of onset latency in the ventral pathway compared to the dorsal pathway this will be difficult to achieve.

For reasons like this, Gray (1999) and Von der Malsburg (1999) restrict the role of synchrony of activation to pre-attentive processes in the lower areas of the visual cortex. For similar reasons, Singer (1999) argues that synchrony of activation is mediated mainly by the magnocellular pathway, because information is transmitted much faster in this pathway, and with a high temporal resolution (e.g., Schmolesky et al., 1998). The magnocellular pathway processes in particular motion information (in the dorsal pathway) and large-scale form information (dorsal and ventral pathway), which coincides with the fact that synchrony of activation is most easily found with moving stimuli and visual gratings (e.g., Gawne, 1999; Friedman-Hill, Maldonado and Gray, 2000).

As discussed earlier, Tononi et al. (1992) presented a computational model of the visual cortex in which form, color, motion and location information is processed in separate but interconnected pathways and areas. The model achieved synchrony of activation for the representations of features belonging to the same object. However, in their model, Tononi et al. assumed equal activation delays between any two areas of the visual cortex. But if different activation delays are incorporated in a computational model, loss of synchrony of activation occurs as well (e.g., Ernst, Pawelzik and Geisel, 1995; Lumer, Edelman and Tononi, 1997; Cambell, Wang and Jayaprakash, 1999).

#### 4. Conclusions

We analyzed the role of synchrony of activation in (higher-level) cognitive processes. In particular, we analyzed the question of whether synchrony of activation provides a mechanism for compositional representation in neural systems. Compositional forms of representation are important for higher-level cognitive processes because they provide the systematicity and productivity needed in these processes. Thus, if synchrony of activation could be used to implement compositional representations in neural systems, synchrony of activation would be an important mechanism for the implementation of cognitive processes in the brain.

On the basis of our analysis, we conclude that synchrony of activation does not provide a mechanism for a neural implementation of compositional representation. In each of the models we analyzed, one can indeed identify a level in which only constituents and their relations are represented, as required for a compositional form of representation. However, selective behavior, such as answers to binding questions, is not produced on this level of representation. Instead, the production of behavior in these models depends on the detection of the synchrony of activation

by means of coincidence detectors. Because coincidence detectors are conjunctive representations, they stand in contrast with compositional representations. Hence, the production of behavior does not result from compositional forms of representation in systems that rely on synchrony of activation. In fact, the role of synchrony of activation in the models we analyzed is not to generate compositional representations, but to make effective use of conjunctive representations. Without synchrony of activation, spurious activations of conjunctive representations (a superposition catastrophe) would occur in these models, representing objects or propositions not presented to the model.

Furthermore, misrepresentations will arise when synchrony of activation is used to represent binary or higher order predicates, such as spatial relations or propositions. These misrepresentations (illustrated in Figure 3) are a result of the inability to distinguish between type and token representations by means of synchrony of activation, which is a direct consequence of the inability to represent compositional structures by means of synchrony of activation.

The role of synchrony of activation will be restricted to processes in which information is represented with conjunctive forms of representation in terms of coincidence detectors, and in which large differences in activation delay do not occur. Examples may be found in the pre-attentive grouping processes in the lower areas of the visual cortex, initiated by the magnocellular pathway.

### Acknowledgements

We wish to thank three anonymous reviewers for their comments on the first draft of the manuscript.

### References

- Barlow, H.B., 1972: Single units and sensation: A neuron doctrine for perceptual psychology? *Perceptron* **1**, 371–394.
- Biederman, I., 1987: Recognition by components: A theory of human image understanding. *Psychological Review* **94**, 115–147.
- Bienenstock, E. and Geman, S., 1995: Compositionality in neural systems. In M.A. Arbib, G. Adelman and P. H. Arbib (eds), *The Handbook of Brain Theory and Neural Networks*, Cambridge, MA: MIT Press, pp. 223–226.
- Cambell, S.R., Wang, D.L. and Jayaprakash, C., 1999: Synchrony and desynchrony in integrate-and-fire oscillators. *Neural Computation* **11**, 1595–1619.
- Dennett, D.C., 1991: *Consciousness Explained*. London: Allen Lane.
- Ernst, U., Pawelzik, K. and Geisel, T., 1995: Synchronization induced by temporal delays in pulse-coupled oscillators. *Physical Review Letters* **74**, 1570–1573.
- Farah, M., Humphreys, G.W. and Rodman, H.R., 1999: Object and face recognition. In M.J. Zigmond, F.E. Bloom, S.C. Landis, J.L. Roberts and L.R. Squire (eds), *Fundamental Neuroscience*, San Diego: Academic Press, pp. 1339–1361.
- Friedman-Hill, S., Maldonado, P.E. and Gray, C.M., 2000: Dynamics of striate cortical activity in the alert Macaque: I. Incidence and stimulus-dependence of gamma-band neuronal oscillations. *Cerebral Cortex* **10**, 1105–1116.

- Fodor, J.A. and Pylyshyn, Z.W., 1988: Connectionism and cognitive architecture: A critical analysis, *Cognition* **28**, 3–71.
- Ghose, G.M. and Maunsell, J., 1999: Specialized representations in the visual cortex: A role for binding? *Neuron* **24**, 79–85.
- Gray, C.M., 1999: The temporal correlation hypothesis of visual feature integration: Still alive and well, *Neuron* **24**, 31–47.
- Hummel, J.E. and Biederman, I., 1992: Dynamic binding in a neural network for shape recognition, *Psychological Review* **99**, 480–517.
- Hummel, J.E. and Holyoak, K.J., 1997: Distributed representations of structure: A theory of analogical access and mapping, *Psychological Review* **104**, 427–466.
- Lumer, D.E., Edelman, G.M. and Tononi, G., 1997: Neural dynamics in a model of the thalamocortical system. I. Layers, loops and the emergence of fast synchronous rhythms, *Cerebral Cortex* **7**, 207–222.
- Mehta, S.D., Ulbert, I. and Schroeder, C.E., 2000: Intermodal selective attention in monkeys. I: Distribution and timing of effects across visual areas, *Cerebral Cortex* **10**, 343–358.
- Mountcastle, V.B., 1998: *The Cerebral Cortex*. Cambridge, MA: Harvard University Press.
- Schroeder, C.E., Mehta, S.D. and Givre, S.J., 1998: A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque, *Cerebral Cortex* **8**, 575–592.
- Schmolsky, M.T., Wang, Y., Hanes, D.P., Thompson, K.G., Leutgeb, S. Schall, J.D. and Leventhal, A.D., 1998: Signal timing across the macaque visual system, *Journal of Neurophysiology* **79**, 3272–3278.
- Shadlen, M.N. and Movshon, J.A., 1999: Synchrony unbound: A critical evaluation of the temporal binding hypothesis, *Neuron* **24**, 67–77.
- Singer, W., 1999: Neuronal synchrony: A versatile code for the definition of relations? *Neuron* **24**, 49–65.
- Singer, W. and Gray, C.M., 1995: Visual feature integration and the temporal correlation hypothesis, *Annual Reviews Neuroscience* **18**, 555–586.
- Singer, W., Engel, A.K., Kreiter, A.K., Munk, M.H.J., Neuenschwander, S. and Roelfsema, P.R., 1997: Neuronal assemblies: Necessity, signature and detectability, *Trends in Cognitive Sciences* **1**, 252–261.
- 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**, 417–494.
- Tononi, G., Sporns, O. and Edelman, G.M., 1992: Reentry and the problem of integrating multiple cortical areas: Simulation of dynamic integration in the visual system, *Cerebral Cortex* **2**, 310–335.
- Van der Velde, F., 1997: On the use of computation in modelling behaviour, *Network: Computation in Neural Systems* **8**, 1–32.
- Von der Malsburg, C., 1987: Synaptic plasticity as basis of brain organization. In J.P. Changeux and M. Konishi (eds), *The Neural and Molecular Bases of Learning*, Chichester: John Wiley, pp. 411–431.
- Von der Malsburg, C., 1999: The what and why of binding: The modeler's perspective, *Neuron* **24**, 95–104.