

Prototypical Relations for Cortex-Inspired Semantic Representations

Florian Röhrbein (florian.roehrbein@honda-ri.de)

Julian Eggert (julian.eggert@honda-ri.de)

Edgar Körner (edgar.koerner@honda-ri.de)

HONDA Research Institute Europe GmbH, Carl-Legien-Str. 30
63071 Offenbach am Main, Germany

Abstract

Cognitive systems for the representation of declarative knowledge like semantic networks and other graph-based systems are widely unrelated to characteristic neurobiological mechanisms in the brain. In this contribution we report on our efforts in bridging the gap between typical semantic relations like “is part of”, “has property” etc. and the laminar wiring pattern of the neocortex. Central to our approach is the identification of the cortical column as a basic building block within the relational network. These columns are typically sectioned into subsystems which comprise different horizontal layers and thereby provide different links for forward, backward and lateral processing. We show how these inter-columnar connections can be related to semantic links, which reflect hierarchical knowledge, temporal ordering and ontological relationship. These dimensions are of outstanding interest for most cognitive tasks. But also arbitrary n-ary relationships can be build by representing the relations as nodes and using only the proposed basic link types. As inference mechanism, a simple locally controlled activation spread was applied. It results directly from the intra-columnar connectivity which is uniform for all nodes. We tested the system with large commonsense databases and obtained promising results including predictions, context influences and feature inheritance.

Keywords: cortical column; knowledge representation; relational structures

Introduction

For the representation of relational knowledge in a graph-based model we have developed a neural-symbolic network which combines ideas from classical semantic networks and recent findings of the neocortical wiring. It consists of columnar-like nodes as uniform entities for the representation of all concepts of the domain, including sensory measurements, motor actions, instances and categories. The nodes are connected by a set of directed links, which can be related to columnar subsystem, as we will argue in the next section.

Semantic relations and columnar connections

The biological entity, which in our approach corresponds to a network node, is the cortical column. The column is well known as the basic computational unit in the brain and its six-layered architecture has been addressed by several researchers to unravel the functional role (Raizada & Grossberg, 2003; Lücke & von der Malsburg, 2004; Kupper *et al.*, 2006). Here we concentrate on a network build out of

columnar-like nodes and do not target at a biologically detailed modeling of the single cortical column. The columns are typically sectioned into subsystems (see Fig. 1) which comprise different horizontal layers and thereby provide different links for bottom-up (BU), top-down (TD) and lateral processing.

We refer to a schema described in (Körner, Tsujino & Masutaki, 1997) which assumes six distinct systems, which we will only briefly sketch here: Subsystem A1 receives input from lower cortical areas, subsystems A2 and B2 project to areas higher in the cortical hierarchy, thus establishing together a bottom-up processing stream. Top-down processing is realized via subsystem C2, which projects to lower areas, targeting in cortical layer I (since there are no neurons in this layer, it is not called a subsystem). The two remaining systems are for lateral processing (B1), which comprises many different cell types and can be subdivided further, and a system which sends primarily motor information to subcortical structures (C1).

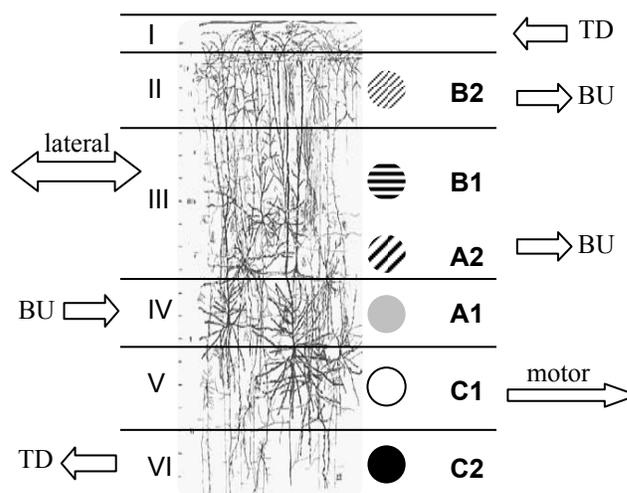


Figure 1: Sketch of the major pathways connecting cortical columns. Shown are cytoarchitecturally defined cortical layers (I-VI, left) and proposed functional subunits (A1 etc., right) with shadings referenced in Fig. 7.

The relevant question in this context now is how semantic links can be ascribed to these pathways which originate from distinct subsystems. A good point to start with will be to look at those semantic relations which seem to be of

ubiquitous importance. Indeed, there seem to be very few basic relations which are relevant for concepts on all layers of abstraction, independent of the actual knowledge domain and these might be grouped according the three dimension of hierarchy, sequence and relationship.

Hierarchies are used all over the neocortex as the core organization principle to deal with the nested structure of the surrounding world. Along this dimension of knowledge chunks the notions of BU and TD processing apply. Knowledge about hierarchical relationships is usually expressed in meronymies and holonymies, but also in relations like “is located in” or in the temporal domain (“happens during” etc). In our system three link types are used to build the chunking hierarchy and, following the basic cortical processing streams, columnar subsystems are assigned to each of them (Fig. 1, for details see e.g. Thomson & Bannister, 2003): A “has component” link, which originates in C2 and projects to layer I of nodes on a lower level (top-down). Two “is component of” links stem from different cortical layers (A2, B2) but terminate both in input layer A1. Together they serve for bottom-up information flow, and just differ in the granularity of transmitted information. Note, that an increased level of detail leads to a hierarchy, in which subclasses are represented above superclasses and instances are represented above categories (see e.g. Quiñ Quiroga *et al.*, 2005), generating a reversed ontological hierarchy.

Sequential information is essential, especially for prediction. We associate corresponding semantic links with the columnar subsystem C1 (compare Lomber & Payne, 2000), but will not make use of it in the work presented here. Instead, we concentrate on ontological knowledge which is expressed in hyponyms and hypernyms. For this dimension (coined “relationship” above) six link types are used: has property / is property of, has subclass / is subclass of and has role / is role of. A suitable columnar subsystem for these connections seems to be B1 because of the existence of distinct functional subsystems within upper layer III (Yoshimura, Dantzker & Callaway, 2005) and the characteristic dense wiring pattern with horizontal connections of different ranges (e.g. Hirsch & Gilbert, 1991). In this line, links denoting subclass relationships connect columns within one level (e.g. within one cortical area), whereas property and role links make inter-area connections, since they connect conceptual representations with more perceptually based ones. Summarizing, we have the following set of link types

- has component / is component of
- has consequence / is consequence of
- has property / is property of
- has subclass / is subclass of
- has role / is role of

All network links proposed here differ in two important aspects to common semantic network links: First, we only use a very restricted set of basic link types, which are biologically justified, since they can be associated with specific neuronal source and target populations each within a specific columnar layer. Second, these links do not vary

from node to node, but are common to all nodes. Not all links, of course, are used by every node, but there are no links which are available only for certain nodes. The motivation for this homogenous layout is that the basic structure of the biological column is widely independent of the cortical site.

In the following the focus will be only on the lateral connections originating and targeting in subpopulations of B1. For details on link types associated with the other subsystems for semantic relations about temporal and spatial ordering, see (Röhrbein, Eggert & Körner, 2007). To motivate the link types associated with B1, we start with quite general considerations concerning the coding of relational structures.

Representation of arbitrary relations

How can arbitrary relations be expressed within a graph-based framework? The common way to represent facts like “John loves Mary” is to have nodes for the concepts involved (“John”, “Mary”) and a directed link between them. The link is typically labeled with the relation, which holds between the connecting concepts (“loves”). Unfortunately, this works only for dyadic relations: As soon as a third concept is involved, like in “Mary gives John a cookie”, the scheme has to be revised.

One option here is to extend the directed links towards “hyperarcs” (see e.g. Harel, 1988). These are either conceptualized as n-ended arcs and then solve only half of the problem, because still only 2 roles are possible. Cases where directed hyperarcs are sufficient are quite restricted, e.g. for relations like lies-between(X,Y,Z,...) which can be represented by setting head H={X} and tail T={Y,Z,...}. Others, e.g. Boley (1992) in his “directed recursive labelnode hypergraphs” propose links which start at the relation node, cut n-1 argument nodes and end at the nth node of the n-ary relation. He criticizes approaches which introduce additional nodes, because they add “pseudo-entities”, but this holds only as long if they cannot be given a reasonable interpretation.

Another solution to handle relations with arity greater than 2 is to have an extra node pointing the all involved concepts (e.g. used in SNePS). The links have to be labeled, otherwise one could not differentiate between the statements “Mary gives John a cookie” and “John gives Mary a cookie”. But here the labels do not reflect the type of relation as for binary relations; rather they specify roles like “giver”, “recipient” and “object”. More as a side effect, in doing so the “artificial” node becomes to represent the whole relation. The provision of labels is a general problem. They are unsatisfactory for several reasons, most importantly for us because biology does not allow for arbitrary link types.

A variant of this approach would be to chunk information into a node which then comes to represent a part of the whole statement like “Mary gives John”. This is useful for computational reasons, since reasoning with nodes of arity beyond 3 have proven to be intractable. Conceptually there

is no further gain in representing parts of a statement, since it makes no sense to dispense completely with separate concepts for both “Mary” and “John”, and the semantic of the link becomes even more obscure.

For a solution without arbitrary link types, two aspects of a “standard link” have to be considered: First, the link has only two ends, and second, there are only two different “values” for the endpoints: “arrow tail” and “arrowhead”, usually associated with an “ingoing” and “outgoing” semantic. For a true extension therefore, a graphical notation is needed based on that depicted in Fig. 2 (b) for a ternary relation. Here the link is allowed to have more than two connecting points and at the same time more than two possible values. For a biological interpretation such a graph-based approach still causes a problem, since there are no different connection endings for neurons that can be associated with arbitrary roles. (In fact, there seem to be different link types and associated with them, different roles, but these are not arbitrarily definable, see above.)

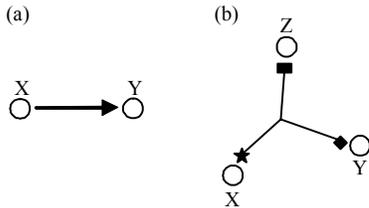


Figure 2: Graphical representations for binary relations (a) have to be extended in two ways to deal with n-ary relations: The number and the type of terminal points. Sketched in (b) is a graphical notation for n=3.

Here we present a very straight-forward solution to this by proposing an additional node for each role. A ternary relation is now represented as sketched in Fig. 3 (b) with new intermediate nodes labeled with digits. This schema can easily be extended for relation with greater arity (c) and is also valid for binary relations (b), thus avoiding any discontinuity. The new nodes have a clear interpretation: For the statement “Mary gives John a cookie” with X=Mary, Y=John and Z=cookie, node 1 represents “Mary acting as giver”, node 2 “John acting as receiver” and node 3 “cookie as a gift”. These nodes can participate in other relational statements, e.g. the node 2, if it is to be expressed that Mary gives some things to some other people. Of course, the same holds for concepts, since these usually participate in different situations having different roles (see simple example below). Fig. 3 (d) shows how the fact “a can is made of aluminum” is represented and indicates the embedding in our columnar framework with links of type has-role / is-role-of (depicted now as solid bidirectional links).

On the conceptual side, the advantage of the proposed schema lies in the uniform treatment of arbitrary n-ary relations. This is opposed to standard semantic networks which show a tendency to binarize not only relations with more than two roles, but also monadic relations like “has

property”. E.g., typical KL-ONE representations are restricted to unary and binary predicates. This is not an inherent restriction and n-ary description logics have been proposed (e.g. Schmolze, 1989), but nevertheless n-tuples for n>2 are usually represented indirectly by reification. There are also technical advantages, since the modeling of relations as nodes allow for inheritance, activation spread etc., which we will elaborate on shortly in the next section.

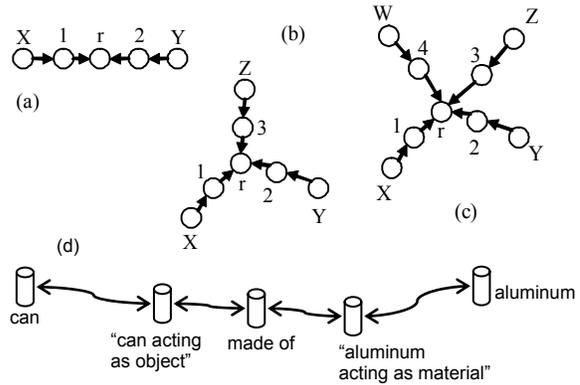


Figure 3: Proposed pattern for representing n-ary relations (a-c, n=2,3,4). The dyadic relation in (d) is represented with 5 nodes connected with has-role / is-role-of links.

Relational prototypes

For every relation which can be expressed in our framework (like “gives”, “made of” etc.) there is a so-called relational prototype, which functions as a template and which is connected with all relations of that type. An example is given in Fig. 4 for the relation “made of”. There are two instances of this relation involving three concepts in the same manner as in Fig. 3 (d), but additionally nodes which code the relation are connected (dashed arrows) with corresponding nodes of the relational prototype. These links are of type has-subclass / subclass-of and thus allow for the inheritance of properties (not included in the figure). The nodes constituting the relational prototype (within the gray oval) get their meaning via their connection to all instances of these nodes.

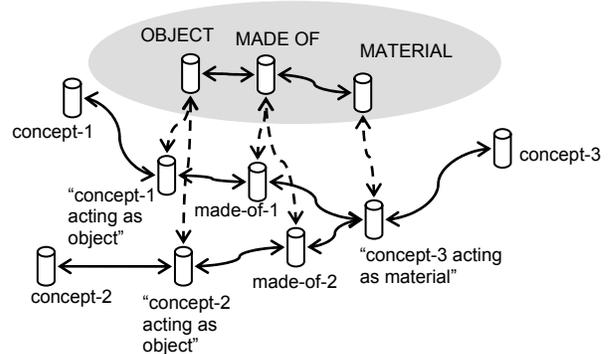


Figure 4: Relational prototypes (grey oval) are component-wise connected with all instances thereof via has-subclass / is-subclass-of links (dashed arrows).

Note that also the cortex seems to build separate representations for different tasks. For spatial cognition tasks e.g. knowledge about spatial relations has to be provided and all the nodes representing instances of these relations should be arranged in neighbored representations. There is also a psycholinguistic justification of treating relations as abstract concepts, which comes from work of Chaffin and Herrmann (1988). They found basic characteristics, which are known to hold for objects, also for relations. These include decomposability, typicality, codability, multiple inheritance and compositionality.

Activation spread

The activation spread results from intra- and inter-columnar connectivity patterns. Internally each node has an activity vector with one entry for each subsystem. The connections between nodes depend on the represented knowledge and follow the rules outlined above. In the current version of the system, we use the simplest rule for intercolumnar connections without weighting and thresholds: The incoming activation a_j of a specific subsystem x_{in} equals the sum of the activations of the corresponding outgoing subsystems of all those nodes i , which are connected to node j :

$$a_j(x_{in}) = \sum_i w_{ji} a_i(x_{out})$$

The processing within a node depends on the intra-columnar wiring, which is handcrafted, but identical for all nodes in the network. We omit here all definitions except those which are made use of in the example which follows in the next section, i.e. we focus on the subsystems marked in grey in Fig. 5.

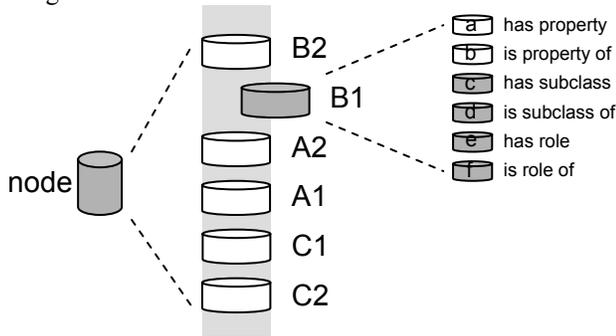


Figure 5: Columnar network nodes are subdivided into functional subsystems and associated semantic link types. In this report the focus is mainly on the shaded parts of B1.

The intra-columnar propagation rules for the activities $a_j(x_{out})$ of a node j are defined in the following for the subsystems x_{out} of B1. We use the abbreviations B1a, ..., B1f given in Fig. 5.

has subclass:

$$a_j(B1c_{out}) = a_j(B1b_{in}) + a_j(A1) + a_j(B1c_{in}) + 2 * a_j(B1f_{in})$$

is subclass of:

$$a_j(B1d_{out}) = a_j(B1a_{in}) + a_j(layerI) + a_j(B1d_{in})$$

has role:

$$a_j(B1e_{out}) = 0.5 * a_j(layerI) + a_j(B1e_{in})$$

is role of:

$$a_j(B1f_{out}) = a_j(B1e_{in}) + 3 * a_j(B1c_{in}) * a_j(B1f_{in})$$

Note, that in all cases the activity vector remains unchanged, unless the incoming activity changes, i.e. there is no automatic fading away.

Toy example

Network

We tested our network with large knowledge databases, i.e., all relations are extensionally defined. In the following we demonstrate the behavior with a toy example consisting of four pieces of relational knowledge, which are fed into the system:

```
can is-made-of aluminum
can is-used-for drinking
can is-used-for gaming
car is-used-for driving
```

A representation of these statements involves six concepts (for can, car, aluminum etc.) and two relational prototypes (is-made-of and is-used-for). All nodes and relations necessary for representing the knowledge are generated automatically resulting in the network shown in Fig. 6.

This example requires the use of only two different link types: has-role / is-role-of links (depicted as solid bidirectional arrows) and has-subclass / is-subclass-of links (dashed bidirectional arrows). Note that role nodes can be part of more than one relational statement (here node “tool-can”).

Task

Let's assume that an object has been presented to the system, which was identified as a can. This successful recognition leads to an activation of the column representing “can”. In the next time step one might wish to ask the system about the usage of this object. This situation can be directly expressed in our network through the activation of two nodes: Node “can” receives top-down input via layer I which activates neurons in subsystem C2 and parts of subsystem B1. We choose an activity value of 1 and set:

$$a_{can}(layerI) = 1$$

Node “usage” is activated bottom-up leading to a firing of neurons in subsystem A1. For simplicity, the same activity value is assumed here:

$$a_{usage}(A1) = 1$$

Starting at these two nodes, the activation spreads according to the proposed intra-columnar wiring and according to the connections between nodes specified by entries of the knowledge base.

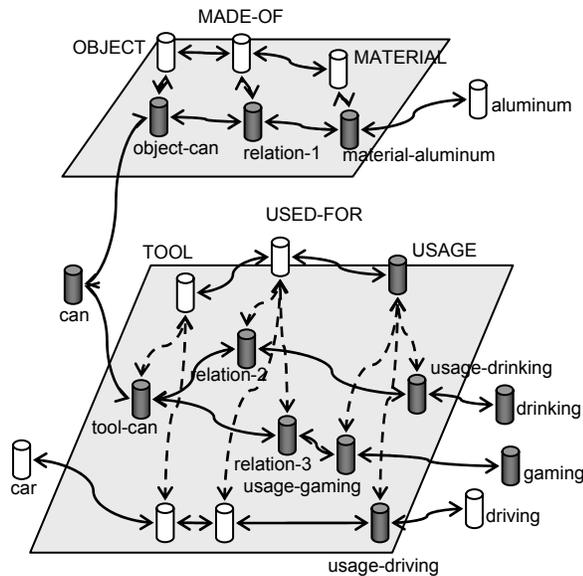


Figure 6: Example network with representations for relations made-of and used-for. See text for details.

Result

In Fig. 6 all nodes which receive activation are colored: The nodes which received input, several nodes coding the relation and two nodes which represent a concept: “drinking” and “gaming”. This highly selective activation of relevant nodes becomes even more important if we consider realistic knowledge networks like that one we obtained by using freely-available ontological and commonsense databases (see Röhrbein et al., 2007). They typically comprise hundreds of instances for one relational prototype, but the only relational structures which get completely activated are those matching with the concepts contained in the “question”, in this example “can” and “USAGE”. As can be seen from the propagation rules defined above, this is due to the nonlinearity for the is-role-of activity, which leads to the desired gating behavior.

For a quantitative comparison the contribution of the different subsystems to the total node activity can be found in Fig. 7. The depicted activity vectors of all involved nodes show the highest activation for nodes “drinking” and “gaming”. Clearly, the gained sum activity scales with the provided input values, which were here set to 1, but the different subsystems’ contributions depend on the

weightings in the definitions above. Moreover, it is not quite clear what should be taken as the “sum activity” of a column (see e.g. measurements by Staiger et al., 2000). On the other hand, the provision of an “answer” as system output, which might consist of a short list of top-ranked nodes, is considered rather as a side effect. The more important result in our view lies in the selective activation of relevant substructures which can be used for subsequent processing steps.

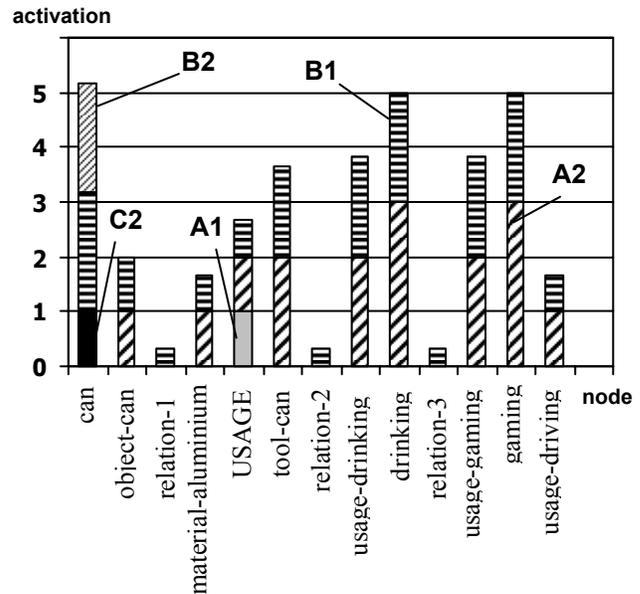


Figure 7: Activity vectors resulting from C2 activation of “can” and A1 activation of “USAGE”. Shading styles refer to columnar subsystems in Fig. 1.

Discussion

A similar node-based representation is used in LISA (e.g. Hummel & Holyoak, 2003), in which a form of symbolic connectionism is proposed which also avoids labeled links for the representation of relational structures. Hummel and Holyoak argue for a 4-tiered hierarchical schema comprising “propositional units”, “role-binding units”, “token units” and “semantic units”. These nodes roughly correspond to neurons, whereas in our approach only one uniform unit is assumed which is related to a larger biological entity, the cortical column. This enables us to differentiate e.g. between superclasses and properties which are treated uniformly in LISA as features at the level of “semantic units”. Related constructs can also be found in connectionist modeling approaches of linguistic aspects. E.g., Hadley and Cardei (1999) introduced p-nodes, which are clusters consisting of a core node connecting sequence nodes with role nodes. Quite similarly to our approach, these role nodes are for the binding of concepts to appropriate roles, but here the possible roles are restricted to a predefined and fixed set of only three role nodes: concepts can be linked either to an

“agent role”, an “action role” or a “patient role”. In SHRUTI (Shastri, 1999), “focal-clusters” represent n-ary relations and contain beside n role nodes also a number of special-purpose nodes like enablers and collectors. Shastri does not relate this unit to biological mechanisms like the cortical column, but it might be worthwhile to pursue that direction.

In general, the coding of higher-valence relations by introducing additional nodes has already been recommended by Levesque, Brachman (1984) and is proposed also in recent approaches (e.g. Schultheis, Barkowsky & Bertel, 2006), but without considering the need for having role nodes. Another example would be the “relational element theory” put forward by Chaffin and Herrmann (1988). In their investigation on analytical vs. unitary approaches to semantic relations they also propose a decomposition into relational elements (like “agent” and “instrument”), but these are more like properties of relations, e.g. “dimension”, “discrete” etc. than role nodes proper. Furthermore they consider only binary relations.

A final note should be made with respect to the postulation of Firstness, Secondness and Thirdness as the fundamental ontological entities (see Sowa, 2000). This trichotomy has not always been interpreted uniformly, but it seems to fit to our graphical representation, where 1st-ness corresponds to the objects per se, 2nd-ness to the properties of an objects and 3rd-ness to the relation between objects (see Fig. 8).

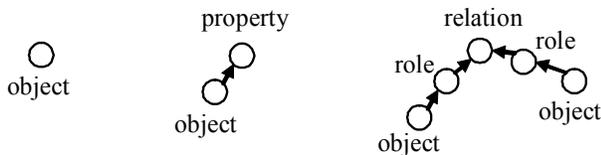


Figure 8: Peircean notions of Firstness, Secondness and Thirdness as major ontological distinctions might be related to network patterns which emerge from our approach.

References

- Boley, H. (1992). Declarative Operations on Nets. In Fritz Lehmann, editor, *Semantic Networks in Artificial Intelligence*, 23, 601-637.
- Chaffin, R., & Herrmann, D.J. (1988). The nature of semantic relations: a comparison of two approaches. In: Evens, M.W. (ed.) *Relational models of the lexicon: Representing knowledge in semantic networks*. Cambridge University Press, 289-334.
- Hadley, R.F. & Cardei, V.C. (1999). Language acquisition from sparse input without error feedback. *Neural Networks*, 12, 217-235.
- Harel, D. (1988). On visual formalisms. *Communications of the ACM*, 31, 514-530.
- Hirsch, J.A. & Gilbert, C.D. (1991). Synaptic physiology of horizontal connections in the cat's visual cortex. *The Journal of Neuroscience*, 11, 1800-1809.
- Hummel, J.E. & Holyoak, K.J. (2003). A Symbolic-Connectionist Theory of Relational Inference and Generalization. *Psychological Review*, 110, 220-264.
- Körner, E., Tsujino, H. & Masutani, T. (1997). A cortical-type modular neural network for hypothetical reasoning. *Neural Networks*, 10, 791-814.
- Kupper, R., Knoblauch, A., Gewaltig, M.-C., Körner, U. & Körner, E. (2006). Simulations of signal flow in a functional model of the cortical column. *Neurocomputing*, doi:10.1016/j.neucom.2006.10.085.
- Levesque, H.J. & Brachman, R.J. (1984). A fundamental tradeoff in knowledge representation languages. In Brachman, R.J. & Levesque, H.J. (eds.), *Readings in Knowledge Representation*. Morgan Kaufmann, Los Altos, California, 41-70.
- Lomber, S.G. & Payne, B.R. (2000) Translaminar differentiation of visually guided behaviors revealed by restricted cerebral cooling deactivation. *Cerebral Cortex*, 10, 1066-1077.
- Lücke, J & von der Malsburg, C. (2004). Rapid Processing and Unsupervised Learning in a Model of the Cortical Macrocolum. *Neural Computation*, 16, 501-533.
- Quian Quiroga, R., Reddy, L., Kreiman, G., Koch, C. & Fried, I (2005). Invariant visual representation by single-neurons in the human brain. *Nature*, 435, 1102-1107.
- Raizada, R. & Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: Computational clues from the visual system. *Cerebral Cortex*, 13, 100-113.
- Röhrbein, F., Eggert, J. & Körner, E. (2007). A Cortex-Inspired Neural-Symbolic Network for Knowledge Representation. In: *Proceedings of the IJCAI-07 Workshop on Neural-Symbolic Learning and Reasoning (NeSy-07)*. CEUR Workshop Proceedings.
- Schmolze, J.G. (1989). Terminological Knowledge Representation Systems Supporting N-ary Terms. In: Brachman, R.J., Levesque, H.J., Reiter, R. (Eds.): *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kaufmann: 432-443.
- Schultheis, H., Barkowsky, T. & Bertel, S. (2006). LTM-C: An Improved Long-Term Memory for Cognitive Architectures. *Proceedings of the 7th International Conference on Cognitive Modeling ICCM 2006*: 274-279.
- Shastri, L. (1999). Advances in SHRUTI - A neurally motivated model of relational knowledge representation and rapid inference using temporal synchrony. *Applied Intelligence*, 11, 79-108.
- Sowa, J.F. (2000). *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks Cole Publishing Co., Pacific Grove, CA.
- Staiger, J.F., Kötter, R., Zilles, K., Luhmann, H.J. (2000). Laminar characteristics of functional connectivity in rat barrel cortex revealed by stimulation with caged-glutamate. *Neuroscience Research*, 37, 49-58.
- Thomson, A.M. & Bannister, A.P. (2003). Interlaminar connections in the neocortex. *Cerebral Cortex*, 13, 5-14.
- Yoshimura, Y., Dantzker, J.L. & Callaway, E.M. (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature*, 433, 868-873.