

SPATIAL REASONING AND CONNECTIONIST INFERENCE*

Antje Beringer[†] and Steffen Hölldobler

Intellektik, Informatik, TH Darmstadt

Alexanderstraße 10

D-6100 Darmstadt (Germany)

{antje, steffen}@intellektik.informatik.th-darmstadt.de

Franz Kurfeß[†]

Neuroinformatik, Universität Ulm

Oberer Eselsberg

D-7800 Ulm (Germany)

franz@neuro.informatik.uni-ulm.de

Abstract

The system presented in this paper is based on ideas taken from Automated Reasoning and Connectionism. The connectionist inference system CHCL is applied to solve the question whether a given sentence correctly describes the spatial relations of objects shown in a visual scene. The sentence is posed in relational form to a set of rules defining spatial relations among objects and a set of facts describing the objects and their relations in the visual scene. These facts are abductively inferred from an analog representation of the visual scene.

1 Introduction

Intellectics, ie. Artificial Intelligence and Cognitive Science, is an interdisciplinary field whose goals are to understand and explain intelligence on the one hand and to develop computational models which show intelligent behavior on the other hand [Bibel, 1992]. In recent years the field has developed in many subfields and these subfields are often developed as independent areas. Automated Reasoning and Connectionism are such subfields. Originally, the goal of Automated Reasoning was to formalize human thought. Much work in this area is directed toward the development of automated theorem provers for certain logics. But although these systems show quite impressive computational power, they are often not adequate as it takes them a long time to solve problems which humans solve almost immediately.

Connectionist and PDP models, on the other hand, have recently attracted much attention as they were applied to models of behavior which match psychological data, are biological plausible, and computationally efficient. But considerably little experience has been gathered so far in modelling high-level cognitive tasks such as knowledge representation and reasoning. Connectionist systems are not yet able to handle inference, the main process of cognition [Fodor and Pylyshyn, 1988], and Paul Smolensky [1987] has argued that *connectionist systems may well offer an opportunity to escape the brittleness of symbolic AI systems ... if we can find ways*

of naturally instantiating the power of symbolic computation within fully connectionist systems.

Unfortunately, the relationship between automated reasoning and connectionist models has often been neglected. There are only few exceptions like the work by Gadi Pinkas [1991b; 1991a], who investigated the relationship between propositional reasoning and energy minimization, or J. A. Barnden [1989], who has given a neural network implementation of syllogistic reasoning. L. Shastri and V. Ajjanagadde [1993] even talk about the *AI paradox* which means the gap between the ability of humans to draw a variety of inferences effortlessly, spontaneously, and with remarkable efficiency and the results about the complexity of reasoning. As shown in [Beringer and Hölldobler, 1993] Shastri and Ajjanagadde's connectionist inference system can be simulated by a classical deductive system if appropriate reduction techniques are applied in parallel.

CHCL is another attempt to investigate the relationship between automated reasoning and connectionism [Hölldobler, 1990a; Hölldobler and Kurfeß, 1992]. CHCL is a purely connectionist inference system for Horn logic, based on a connectionist unification algorithm [Hölldobler, 1990b]. It utilizes the connection method [Bibel, 1987], where a proof for a formula consists of a so-called spanning set of simultaneously unifiable pairs of literals (occurring in this formula). CHCL generates sequentially all spanning sets for a given formula and unifies the pairs of literals in these sets in parallel. It also contains various reduction techniques which are applied in parallel and essentially reduce the number of spanning sets.

In this paper we show how a slightly modified version of CHCL can solve spatial reasoning tasks occurring in the miniature language acquisition project L_0 . L_0 is a project to investigate how could we learn to describe what we see [Feldman *et al.*, 1990]. The task is to construct a system that decides after a learning period whether a statement is a true description of a given picture. The pictures show two-dimensional scenes with simple geometric objects. The statements are drawn from a finite corpus capturing (as closely as possible) the meaning of English descriptions. As an example, Figure 1 shows a typical scene with a true description.

A first attempt to give a connectionist semantics for the spatial relations has been made by S. H. Weber [1991]. Since the prohibitive combinatorics makes it impossible to simply tabulate all geometric relationships,

* A long version of this paper is available from the authors.

[†]A. Beringer is supported by DFG, project MPS, grant no. HO 1294/3-1. F. Kurfeß is supported by BMFT, project WINA, contract no. 413-4001-01 IN 103 E/9.




Figure 1: There is a dark object right of the square and a large, light object below the circle.

her idea is to employ an iterative search controller providing a focus on certain candidate objects. An open problem in Weber's approach is the question how such a controller can be modelled in a connectionist setting and this is one of the tasks considered here.

The connectionist system presented in this paper decides whether a statement is a true description of a visual scene. The objects in the scene are represented analogically as a pattern of activation in a matrix of units called plane. As only a limited number of objects is present in each scene, the objects are phase-coded as in [Shastri and Ajjanagadde, 1993]. The external activation of the objects leads to the recruitment of a reference unit for each object such that (after the recruitment) the excitation of the reference unit triggers the activation of the object in the plane and vice versa. The features of each object — ie. its shape, size, and shade — are extracted using a trained multi-layer feed-forward network such that this information becomes available for CHCL. The rules defining spatial relations like $bel(T, L) \Leftarrow rbel(L, R) \wedge in(T, R)$ — expressing that a trajectory T is below a landmark L if there is a region R below L and T is in R — are represented as usual in CHCL¹. The facts for the relations $rbel$ and in are not explicitly represented, rather there is a dummy for the two relations which tells CHCL that there are facts for $rbel$ and in . As described in the next paragraph these facts are inferred abductively as soon as necessary. Finally, the sentence is posed as a query.

As CHCL generates the spanning sets of a formula sequentially, it provides the controller for focussing on certain objects and spatial relations. If, for example, the controller requires to check whether the object b is below the object c as shown in Figure 1, then the dummy for $rbel$ and in is activated with L bound to c and T bound to b . With the help of the reference unit for c , the activation of the dummy will trigger a downward spreading of activation in the plane originating from the units representing c . If this activation reaches the units representing b , then b is in the region below c .

In Section 2 the kind of spatial reasoning problems considered here is presented. Section 3 briefly describes CHCL and the main Section 4 describes the modifications and extensions of CHCL needed to solve the spatial reasoning problems. Finally, Section 5 discusses the results and possible directions for future work.

¹Variable symbols begin with upper case letters, whereas predicate and constant symbols begin with lower case letters. Variables are considered to be universally quantified.

2 Spatial Reasoning Problems in L_0

As mentioned, the problem is to decide whether a given statement is a true description of a given visual scene containing simple graphical objects. We are interested in their shape (square sq , triangle tr , or circle ci), shade (dark da or light li), and size (large l , medium m , or small s). The scene is given as a picture from which we have to extract the features. In the example shown in Figure 1 we should extract the facts $obj(a, sq, da, s)$, $obj(b, sq, li, l)$, and $obj(c, ci, li, s)$, with system generated internal names a , b , and c . The spatial relations are defined by universally closed rules of the general form

$$relation(T, L) \Leftarrow spec_region(L, R) \wedge in(T, R) \quad (1)$$

where $relation$ is abv , bel , lo , or ro denoting that the trajectory T is above, below, left or right of the landmark L , respectively. in denotes that T is in region R and $spec_region$ is $rabv$, $rbel$, rlo , or rro denoting the region R above, below, left or right of L , respectively.² To solve the example shown in Figure 1 we need the rules

$$\begin{aligned} ro(T, L) &\Leftarrow rro(L, R) \wedge in(T, R), \\ bel(T', L') &\Leftarrow rbel(L', R') \wedge in(T', L'). \end{aligned}$$

Such rules are applicable if their conditions are satisfied. This requires to determine regions in the visual scene and to check whether an object is in a certain region. In our example this requires to determine the region r_1 right of object b ($rro(b, r_1)$) and the region r_2 below object c ($rbel(c, r_2)$), and to check whether object a is in r_1 ($in(a, r_1)$) and b is in r_2 ($in(b, r_2)$).

A statement is a conjunction of atomic propositions describing objects and their spatial relations. Statements are quite simple English sentences and we may assume that they are in relational form. The sentence in Figure 1 can thus be expressed as the query³

$$\begin{aligned} \Leftarrow & \quad obj(X, -, da, -) \wedge obj(Y, sq, -, -) \wedge ro(X, Y) \\ & \wedge obj(U, -, li, l) \wedge obj(V, ci, -, -) \wedge bel(U, V). \end{aligned}$$

Logically, to determine whether a statement correctly describes a picture we must decide whether the query representing the statement is a logical consequence of the facts about objects and regions extracted from the picture and the set of rules defining the spatial relations.

3 CHCL

CHCL is a connectionist inference system for Horn logic [Hölldobler, 1990a; Hölldobler and Kurfeß, 1992]. From its expressive power it can easily handle the kind of problems presented in Section 2 if all the rules and facts are given. CHCL is based on the connection method [Bibel, 1987]. This method tells us that a proof of a formula is found, if there exists a spanning set of pairs

²One should observe that Shastri and Ajjanagadde's [1993] as well as Lange and Dyer's [1989] inference system cannot handle this kind of rules as the variable Z does not occur in the conclusion of the rule but twice in its conditions.

³The symbols “-” denote different variables, where we do not care about the bindings. The connectionist inference systems proposed in [Shastri and Ajjanagadde, 1993] and [Lange and Dyer, 1989] again cannot handle this kind of queries as unbound variables like the X occur more than once.




Figure 2: The formula representing the problem depicted in Figure 1, where the logical connectives are omitted. The lines show unifiable pairs of literals. The set of pairs of literals shown by regular lines is a spanning set.

of literals such that all pairs are simultaneously unifiable. A spanning set can informally be defined as follows: Each literal occurring in the query must be in the spanning set and, if the conclusion of a rule is in the spanning set, then each literal occurring in the conditions of the rule must be in the spanning set. It is easy to check that the regular curves in Figure 2 form a spanning set. Moreover, all pairs of literals in this set are simultaneously unifiable yielding the answer substitution $\{X \mapsto a, Y \mapsto b, U \mapsto b, V \mapsto c\}$. For all other spanning sets the pairs of literals are not simultaneously unifiable.

CHCL contains a subnetwork to represent the pairs of literals which are potential elements of spanning sets. For each pair of literals l this subnetwork contains an output unit which becomes active iff l is an element of the current spanning set. The subnetwork resembles a Jordan network [Jordan, 1986] and sequentially generates the various spanning sets of a formula. As soon as a spanning set has been determined, all pairs of literals occurring in this set have to be simultaneously unified. This is done with the help of the connectionist unification algorithm presented in [Hölldobler, 1990b].

Terms are represented in a so-called term layer using a role-filler representation as shown in Figure 3. Each element of the term layer is a selfexcitatory binary threshold unit. The eight units representing shape, shade, and size are referred to as the *feature vector*. The large dark squares ■ represent the terms in the original formula. If the spanning set contains a pair of literals like $\langle \neg obj(X, -, da, -), obj(a, sq, da, s) \rangle$, then the representations of the terms of these literals in the term layer are




Figure 3: The term layer for the example shown in Figure 2. The terms of the literal next to each row or pair of rows are represented by the large dark squares ■.

put on top of each other. This is simultaneously done for all elements of the spanning set. In Figure 3 the medium sized dark squares ■ show the result of this process for the spanning set defined by the regular curves in Figure 2. Now, the terms are unified.

The problems considered in this paper are such that terms are constants and variables only. Hence, unification is successful if multiple occurrences of the same variable are bound to the same term and two different constants do not appear in the same position. The bindings are represented as the activation pattern of the rows in the term layer. For example, from the activation pattern in the first row of the term layer shown in Figure 3 we learn that the variables X and T are bound to the constant a . The unification algorithm further activates units such that whenever a variable occurs in more than one row, then the activation patterns of these rows are recursively put on top of each other until eventually they become identical. This takes no longer than $O(\log(n))$ time in the worst case, where n is the number of variables occurring in the formula. The small dark squares ■ show the result of this process for the running example. Two different constants appear in the same position if in one line more than one unit for the shape, shade, size, or references is active. As this is not the case in Figure 3 all elements of the spanning set are simultaneously unifiable and, hence, a proof is found. Note that an attempt to unify $obj(X, -, da, -)$ and $obj(a, sq, li, l)$ would result in the simultaneous activation of the units representing darkness (da) and lightness (li) for the two literals.

4 Spatial Reasoning with CHCL

A slightly modified version of CHCL can be used to deal with the problems described in Section 2. We consider a fixed query together with a fixed number of copies of each rule. The maximal number of simultaneously considerable objects is restricted to ω , where $\omega = 7$ seems to be a biologically and psychologically plausible restriction [Miller, 1956; Shastri and Ajjanagadde, 1993].

The picture is represented as a pattern of activation in a matrix of units, called *plane*, where each unit represents an atomic part of the picture. To distinguish the objects considered simultaneously, they are coded in different phases. So the matrix units are phase sensitive binary threshold units. Those collectively representing an object are activated in the same unique phase.

The objects and their features must be extracted from the analogue representation. The object descriptions are stored in feature vectors as described in Section 3. To establish the correspondence between the units representing a certain object in the plane and the associated object description in the term layer we use phase sensitive *reference units*, which are recruited in an extra step. A reference unit will be activated if the units of the corresponding object in the plane are active, and vice versa. Thus, by activating the reference unit via the feature vector, the analogue object representation can be retrieved and by activating the reference units via the analogue object representation, i.e. the units in the plane, the object's internal name and features can be recovered.

Spatial relations between objects must be inferred abductively. In particular, we must consider regions in the plane wrt a given landmark and check whether a given trajectory is in such a region. The regions are identified by spreading of activation originating from the units representing the landmark. Therefore, each of the units in the plane is connected to its neighbours and these connections are controlled by phase sensitive *direction gates*. If, for example, the region below a landmark is to be considered, the gates on the connections going downwards from the units in the plane are opened through activation of the corresponding *direction unit*. If the units representing the landmark are activated now, this activation spreads downwards in the plane. Since the involved units are phase sensitive, ω different regions can be generated simultaneously without cross-talk.

After stating the sentence as a query and presenting the picture by externally activating appropriate matrix units the network's tasks are as follows.

1. Recruit a reference unit for each object.
2. Extract the features of the objects.
3. Generate a new spanning set S . If there exists none, stop with NO.
4. Simultaneously unify the pairs of literals occurring in S . If this results in a failure go to 3.
5. Try to abductively infer the regions and related facts needed for S being a proof for the query. If this is possible, then stop with YES; otherwise go to 3.

The recruitment of the reference units can be done as described in [Diederich, 1988] or [Feldman, 1982]. The reference units are connected bidirectionally to every matrix unit via links with initially low weights




Figure 4: The activation of the matrix units for the example plane with the resulting condensed feature vector CFV and the derived feature vectors.

which are adjusted phase dependently by Hebbian learning. The connections from the matrix to the reference units are in addition normalized [Grossberg, 1988] such that a threshold of .7, for example, means that the reference unit should be activated if more than 70% of the units representing the object in the plane are active. In other words, if at least 70 % of an object are in a certain region, the object is said to be in this region.

The feature extraction component mainly consists of a feed forward network built of the matrix units as input units, one hidden layer, and eight phase sensitive units representing a condensed version of the feature vectors (CFV), where the features of all objects are coded in parallel by a phase dependent representation. The correct weights of the connections are learned by back-propagation. Thereafter, the pattern of activation in the plane causes some units of the CFV to be activated. As each reference unit is activated in a unique phase, they can be used to divide the CFV into at most ω separate feature vectors (see Figure 4). The rows of the feature vectors in the term layer additionally contain ω marker units representing the internal names of objects.

The computation of the spanning sets is explained in Section 3. The number of spanning sets is always finite as we allow only a fixed number of copies for each rule. Thus, if none of the spanning sets constitutes a proof, a user may provide additional copies.

The abductive reasoning component is explained with respect to the example shown in Figure 1. From the feature extraction component we obtain the facts shown in Figure 4. Thereafter, CHCL eventually generates the spanning set (2) except that the facts $rro(b, r_1)$, $in(a, r_1)$, $rbel(c, r_2)$, and $in(b, r_2)$ representing the spatial relations are not present yet. After a successful unification of the pairs of literals in the spanning set, it must be checked whether $ro(a, b)$ and $bel(b, c)$ hold. The rows in the term layer representing the conditions of these rules are partially instantiated to $rro(b, R)$, $in(a, R)$ and $rbel(c, R')$, $in(b, R')$. Figure 5 shows the rows for $rbel(c, R')$ and $in(b, R')$.

The literals in the body of a rule share a *rule unit* being activated by the spanning set unit for the head of the rule (see Figure 5). The rule unit is the dummy establishing that the two literals in the body share a variable not occurring in the head of the rule. Moreover, it repre-




Figure 5: The spatial inference component to infer $bel(b, c)$. Not all links are shown. Everything apart from the "down" unit, the reference units, and the plane is used for exactly one copy of the rule. Phase sensitive binary threshold units are represented by circles labeled with $\sim$. The units of the term layer rows for $rbel$ and in are represented by squares in a row. Modifiers are phase independent inhibitors, represented by little black dots touching the concerned connection. Phase sensitive gates with threshold 1.5 (represented by little circles) control connections within a certain phase. If unspecified, weights equal 1 and thresholds .5.

sents the assumption that facts satisfying the conditions of the rule can be inferred. Figure 5 shows the part of the spatial inference component to infer $bel(b, c)$. The tasks of its units are the following.

The spanning set unit belonging to the conclusion of the rule activates two other units, viz. the binary failure unit and the rule unit.

The failure unit signals that the rule has been used. It is connected to the answer component of CHCL. This connection is inhibited if spatial relations satisfying the conditions of the rule can be inferred. Otherwise, the activation of the failure unit forces a negative answer.

The rule unit must trigger the abductive inference process and, therefore, has three tasks, viz. to open the correct gates in the plane, to activate the units representing the landmark — and, thus, to activate the missing region —, and to check whether the trajectory is in the region. The rule unit is phase sensitive and, hence, ω different spatial relations can be tested simultaneously.

If a rule unit is excited in phase p , it causes the opening of the appropriate direction gates in the plane for p by exciting the corresponding direction unit in phase p . In Figure 5 the down unit is activated. As the marker c in the term layer row for $rbel$ is activated, the rule unit excites the reference unit for c in phase p by opening the corresponding gates. The reference unit excites the units representing c in the plane in phase p . (The connections between the reference units and the plane are not shown in Figure 5.)

To infer whether the trajectory b is in the region below c , the rule unit opens the gates on the connections from marker units of the in row to the gate on the connection from the reference units to the failure unit in phase p . Thus the activation of the marker unit b is propagated in phase p and causes the opening of the gate behind reference unit b in phase p . As b is below c in the plane, the

activation of the corresponding units for c in the plane spreads downwards and reaches the units for b . Thus these units and, consequently, the reference unit for b become active in phase p , too. Reference unit b now inhibits the connection from the failure unit facilitating a positive answer. The marker units for $rbel$ inhibit the connections from the same marker in in , avoiding inferences like $bel(a, a)$.

Finally, as different phases are used for the rule units, the relations $bel(b, c)$ and $ro(a, b)$ can be tested simultaneously. The activation of the units representing object b does not spread to the right, because the gates to the right are opened only in a different phase corresponding to the query $ro(a, b)$.

5 Discussion

We would like to build a system which decides whether a statement in natural language is a true description of a given picture. Such a system must be capable of reasoning, learning, and autonomous operation; it should be massively parallel and physiologically plausible. The main emphasis in this paper is on the reasoning component. Therefore, we assume that the description is available in a logic notation and that the picture is given as a pattern of activation in a matrix of units. In this setting a number of technical problems have to be solved: an adequate internal representation has to be found; the reasoning mechanism has to be powerful enough; the correlation between the objects and their features has to be preserved; the description of features cannot be in absolute terms and linguistic variables [Zadeh, 1989] might be necessary to characterize their relative values; the variable binding problem; the treatment of alternative solutions; abductive reasoning via the generation of required facts on demand; the concordance between the expression of spatial relations in language and the relative positions of objects in the visual scene.

The system presented in this paper achieves most of the goals and solves the technical problems in a prototypical way. The reasoning mechanism is essentially the same as in CHCL. In contrast to other connectionist inference mechanisms (eg. [Shastri and Ajjanagadde, 1993; Lange and Dyer, 1989]) CHCL is powerful enough to deal with the class of problems required here. Our system is capable of learning to associate objects and features. A backpropagation network is trained to extract features from the pixel representation; no information about the relative or absolute position is preserved. Reference units are recruited to relate the objects in the geometric scene to the internal representation of their features. The quest for autonomous operation obviously is relative: after extensive training we expect the system to be capable to handle previously unseen pictures and new sentences. CHCL is a massively parallel inference system in the sense of [Feldman and Ballard, 1982]. Computation is performed by spreading activation.

As for the more technical problems, we have shown that the internal representation as used in CHCL can be applied without major problems to the specific task of spatial reasoning. Physiological and cognitive plausibility can be considered on various levels. On a high level, the treatment of alternative solutions in CHCL is

sequential, which seems to be the same for reasoning tasks which require deliberate thinking in humans. On a lower level, there is no strong evidence that reasoning problems are represented in the brain as matrices mirroring the structure of predicates and terms. As a model, however, neural networks can be represented in a similar way as associative memories [Palm, 1980], which again have a strong similarity to the structure of the matrix used in CHCL. For the bindings between objects and their features [Rock and Palmer, 1990] and for testing spatial relations we use phase-sensitive units similar to the ones used in [Shastri and Ajjanagadde, 1993]. This concept seems to be quite plausible and allows to treat a moderate number of objects and relations simultaneously. The use of linguistic variables is implicit in the learning process for the features and their values. However, it is not yet supported from the language end.

A solution for the variable binding problem as well as for the treatment of alternative solutions is already given through CHCL: the correspondence between variables and their values are represented in the term layer as shown in Figure 3, and alternative solutions correspond to different spanning sets. In contrast to conventional inference mechanisms this representation even allows for a straightforward realization of abductive reasoning: whereas in conventional systems all possible combinations of features and objects would have to be generated and stored as facts, our system initiates the generation of a fact only when it is really needed in the reasoning process. In our present system, the four elementary spatial notions below, above, left-of, and right-of are built into the system via direction units, which control the direction gates for a specific flow of activation (see Figure 5). These notions are used to compose the spatial relations used in the language part.

In total, the system as presented here solves the problem of checking the correspondence between a simple visual scene and a simple natural language characterization of the spatial relations between objects in that scene. With a moderate enhancement of the basic CHCL mechanism, it is possible to extract the necessary information from the analogue representation of the scene, without a necessity to generate an excessive amount of facts describing all its spatial aspects. The system is currently being implemented using ICSIM [Schmidt, 1990].

As for future work many open questions remain: How can the rules be learned? Can we train the system such that spanning sets leading to a proof are generated with a certain priority? What is the meaning of above, below, etc. and when do we call an object small rather than medium, ...?

References

- [Barnden, 1989] J. A. Barnden. Neural-net implementation of complex symbol-processing in a mental model approach to syllogistic reasoning. *Proc. IJCAI*, p. 568-573, 1989.
- [Beringer and Hölldobler, 1993] A. Beringer and S. Hölldobler. On the adequateness of the connection method. *Proc. AAAI*, 1993. (to appear).
- [Bibel, 1987] W. Bibel. *Automated Theorem Proving*. Vieweg Verlag, Braunschweig, second edition, 1987.
- [Bibel, 1992] W. Bibel. *Intellectics. Encyclopedia of Artificial Intelligence*, p. 705-706. John Wiley, NY, 1992.
- [Diederich, 1988] J. Diederich. Connectionist recruitment learning. *Proc. ECAI*, p. 351-356, 1988.
- [Feldman and Ballard, 1982] J. A. Feldman and D. H. Ballard. Connectionist models and their properties. *Cognitive Science*, 6(3):205-254, 1982.
- [Feldman et al., 1990] J. A. Feldman, G. Lakoff, A. Stolcke, and S. H. Weber. Miniature language acquisition: A touchstone for cognitive science. *Proc. COGSCI*, p. 686-693, 1990.
- [Feldman, 1982] J. A. Feldman. Dynamic connections in neural networks. *Biological Cybernetics*, 46:27-39, 1982.
- [Fodor and Pylyshyn, 1988] J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Connections and Symbols*, p. 3-71. MIT Press, 1988.
- [Grossberg, 1988] S. Grossberg. Adaptive pattern classification and universal recoding i. parallel development and coding of neural feature detectors. *Neurocomputing: Foundations of Research*, chapter 19. MIT Press, 1988.
- [Hölldobler and Kurfess, 1992] S. Hölldobler and F. Kurfess. CHCL - A connectionist inference system. *Parallelization in Inference Systems*, p. 318-342. LNAI 590, 1992.
- [Hölldobler, 1990a] S. Hölldobler. CHCL - A connectionist inference system for Horn logic based on the connection method and using limited resources. TR-90-042, ICSI, Berkeley, 1990.
- [Hölldobler, 1990b] S. Hölldobler. A structured connectionist unification algorithm. *Proc. AAAI*, p. 587-593, 1990.
- [Jordan, 1986] M. I. Jordan. Attractor dynamics and parallelism in a connectionist sequential machine. *Proc. COGSCI*, p. 531-546, 1986.
- [Lange and Dyer, 1989] T. E. Lange and M. G. Dyer. High-level inferencing in a connectionist network. *Connection Science*, 1:181-217, 1989.
- [Miller, 1956] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63(2):81-97, 1956.
- [Palm, 1980] G. Palm. On associative memory. *Biological Cybernetics*, 36:19-31, 1980.
- [Pinkas, 1991a] G. Pinkas. Propositional non-monotonic reasoning and inconsistency in symmetrical neural networks. *Proc. IJCAI*, p. 525-530, 1991.
- [Pinkas, 1991b] G. Pinkas. Symmetric neural networks and logic satisfiability. *Neural Computation*, 3, 1991.
- [Rock and Palmer, 1990] I. Rock and S. Palmer. The legacy of Gestalt psychology. *Scientific American*, 263:48-61, 1990.
- [Schmidt, 1990] H. W. Schmidt. ICSIM: Initial design of an object-oriented net simulator. TR-90-055, ICSI, Berkeley, 1990.
- [Shastri and Ajjanagadde, 1993] L. Shastri and V. Ajjanagadde. From associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioural and Brain Sciences*, 1993. (to appear).
- [Smolensky, 1987] P. Smolensky. On variable binding and the representation of symbolic structures in connectionist systems. CU-CS-355-87, Univ. of Colorado, Boulder, 1987.
- [Weber, 1991] S. H. Weber. A connectionist semantics for spatial relations. ICSI, Berkeley, 1991.
- [Zadeh, 1989] L. A. Zadeh. Knowledge representation in fuzzy logic. *IEEE Transactions in Knowledge and Data Engineering*, 1, 1989.