# The Analysis of Natural Language Concepts with Connectionist/Symbolic Techniques

Stefan Wermter

1989

COINS Technical Report 89-117

Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003
USA

## Abstract

This technical report describes three approaches for understanding natural language concepts. The first approach integrates semantic and syntactic constraints for structural noun phrase disambiguation. Semantic constraints are learned as distributed representations in backpropagation networks. These semantic constraints are integrated with syntactic constraints using localist relaxation networks.

The second approach describes a model for learning semantic relationships in compound nouns. A connectionist architecture consisting of modular backpropagation networks can learn and generalize basic semantic relationships in compound nouns.

The third approach is a hybrid model which combines symbolic and connectionist techniques for understanding noun phrases. A distributed connectionist level provides a learned semantic memory model. A localist connectionist level integrates semantic and syntactic constraints. A symbolic level is responsible for restricted syntactic analysis and concept extraction.

All three approaches were tested in the domain of the physical sciences. The noun phrases and compound nouns were taken from the Physical Science Laboratory Corpus. Based on our experiments we conclude that a hybrid connectionist/symbolic approach can be a potentially powerful mechanism for learning, representing, and disambiguating noun phrases in real world domains.

## References

These approaches are/will be published as:

Wermter, S. 1989. Integration of Semantic and Syntactic Constraints for Structural Noun Phrase Disambiguation. Eleventh International Joint Conference on Artificial Intelligence, August 1989, Detroit.

Wermter, S. 1989. Learning Semantic Relationships in Compound Nouns with Connectionist Networks. Eleventh Annual Conference of the Cognitive Science Society, August 1989, Ann Arbor.

Wermter S., Lehnert W.G. 1990. A Hybrid Symbolic/Connectionist Model for Noun Phrase Understanding. Connection Science Vol. 1(3).

# Integration of Semantic and Syntactic Constraints for Structural Noun Phrase Disambiguation*

**Stefan Wermter**
Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003

## Abstract

A fundamental problem in Natural Language Processing is the integration of syntactic and semantic constraints. In this paper we describe a new approach for the integration of syntactic and semantic constraints which takes advantage of a learned memory model. Our model combines localist representations for the integration of constraints and distributed representations for learning semantic constraints. We apply this model to the problem of structural disambiguation of noun phrases and show that a learned connectionist model can scale up the underlying memory of a Natural Language Processing system.

## 1 Introduction

The structural and semantic understanding of noun phrases and prepositional phrases is one of the most important tasks for natural language processing systems. Lately issues of prepositional phrase attachment have been addressed in different systems for sentence understanding (e.g. [Wilks et al. 85], [Schubert 86], [Dahlgren and McDowell 86], [McClelland and Kawamoto 86], [St. John and McClelland 88]). These systems focus on deciding whether a prepositional phrase attaches to a verb phrase or a noun phrase, for instance [Wilks et al. 85]:

The woman wanted the dress on the rack.
The woman positioned the dress on the rack.

In the first example "on the rack" attaches to the noun phrase "the dress", in the second example to the verb "positioned".

All these referenced systems emphasize prepositional phrase attachment in sentences of the form <S><VP><NP><PP>, and concentrate on the attachment of a *single* prepositional phrase based on predictive verbal knowledge. However, attachment decisions for *multiple* prepositional phrases have to rely on syntactic and semantic knowledge associated with nouns and prepositions as well. The importance of this knowledge about nouns and prepositions is very obvious for the attachment decisions in isolated noun phrases, as for example in titles of scientific articles. In this paper we restrict our efforts to prepositional attachment in noun phrases using a corpus of titles and scientific articles from the physical sciences, for instance:

Forces on charged particles of a plasma in a cavity resonator.
Irregularities in the drag effects on sputniks.

We describe a two-level architecture for integrating syntactic and semantic constraints to disambiguate PP-attachment in noun phrases. The bottom level consists of backpropagation networks using distributed representations for the semantic relationships between nouns and prepositions. The backpropagation networks are trained with examples of these prepositional relationships for each preposition, so that the backpropagation networks learn the underlying semantic constraints. The top level consists of a relaxation network using localist representations for the integration of syntactic constraints with the learned semantic constraints. This approach allows the disambiguation of noun phrases which the system has not been trained on.

## 2 Noun Features for Prepositional Relationships

Prepositional relationships depend on domain-specific features of the involved nouns. The noun phrases for our experiments were taken from the NPL corpus [Sparck-Jones 76] which contains article titles for scientific and technical domains. Typical examples in the corpus are:

Pulse techniques for probe measurements in gas discharges.
The influence of the radiation intensity on discharges in the Van-Allen-belt.

For each of the 10 most frequent prepositions in the corpus, 100 noun phrases were extracted which contained the specific preposition. The typical structure of the considered noun phrases is a sequence of up to five noun

*This research was supported by the Advanced Research Projects Agency of the Department of Defense, monitored by the Office of Naval Research under contract #N00014-87-K-0238, the Office of Naval Research, under a University Research Initiative Grant, Contract #N00014-86-K-0764 and NSF Presidential Young Investigators Award NSFIST-8351863.

groups each separated by a preposition. The head noun in the noun group was characterized with semantic features. We found the following 16 features useful as a basic representation for the noun groups in this domain (see Figure 1).

| Features | Examples |
|---|---|
| MEASURING-EVENT | Observation |
| CHANGING-EVENT | Amplification |
| SCIENTIFIC-FIELD | Mechanics |
| PROPERTY | Intensity |
| MECHANISM | Experiment |
| ELECTRIC-OBJECT | Transistor |
| PHYSICAL-OBJECT | Earth |
| RELATION | Cause |
| ORGANIZATION-FORM | Layer |
| GAS | Air |
| SPATIAL-LOCATION | Antarctic |
| TIME | June |
| ENERGY | Radiation |
| MATERIAL | Aluminium |
| ABSTRACT-REPRESENTATION | Note |
| EMPTY | Cavity |

Figure 1: Features of the Nouns and Examples

Most nouns have a clear preference for one of the 16 features, for example "June" for TIME. Although prepositional relationships could be defined with one feature class [Herskovits 86], nouns can have more than one feature, for example "radiation" can be a form of ENERGY, and a CHANGING-EVENT. To account for these multiple features of single nouns each noun is represented as a binary vector of length 16.

## 3 The Structural Disambiguation of Noun Phrases

The disambiguation of noun phrases relies on two types of knowledge: first, semantic, domain-dependent constraints for the plausibility of prepositional relationships and second, syntactic, domain-independent constraints for crossing dependencies and locality.

### 3.1 The Bottom Level: Learning Semantic Constraints with Backpropagation Networks

Semantic constraints based on the plausibility of prepositional relationships determine how different prepositional phrases in a noun phrase can attach to one another. In many systems semantic constraints are formulated as rules (e.g. [Wilks et al. 85] [Dahlgren and McDowell 86]). We describe a different approach for learning the semantic constraints in prepositional relationships.

Learning prepositional relationships for different prepositions is defined as learning to differentiate between plausible prepositional relationships and implausible prepositional relationships. Plausible prepositional relationships are relationships which can be true. For instance, the prepositional relationship "radiation in atmosphere" is plausible. Implausible prepositional relationships are relationships which violate semantic restrictions. For instance, the prepositional relationship "symposium in ionosphere" violates semantic restrictions because meetings are not supposed to take place in the upper atmosphere.

Backpropagation networks are useful to learn the plausibility of prepositional relationships between two nouns and to generalize the regularities for the plausibility of pairs of nouns with which the network has not been trained. We used the backpropagation algorithm as described in [Rumelhart et. al 86]. One backpropagation network is used for representing the prepositional relationships for one preposition. Each network consists of 32 input units, 12 hidden units and one output unit (see figure 2). The input units represent the binary features of the two nouns. The output unit is a real value between 0 and 1 representing the plausibility of the prepositional relationship between two nouns. The hidden units represent the learned mapping between the noun features and the plausibility value.

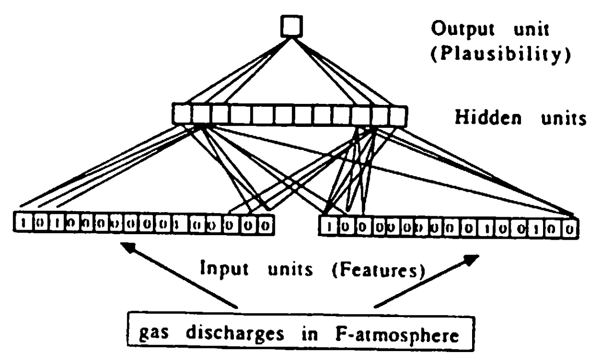


Figure 2: Bottom Level: Backpropagation Network for Learning Prepositional Relationships for the Preposition "in" (only some connections shown)

The backpropagation networks were trained by presenting about 200 training examples for each specific preposition. A training example consisted of the feature representations for the two nouns together with the plausibility value "1" for "plausible" or "0" for "implausible". After the backpropagation networks were trained for 1600 epochs with the training set, each network was tested with the training set and a testing set. The testing set consisted of 30 examples of prepositional relationships (each characterized by 32 features) which the network had not been trained on. A prepositional relationship was considered correct on a scale from 0 to 1 if the value of the output unit was higher than 0.5 for a desired plausible relationship and smaller than 0.5 for a desired implausible relationship. The testing results for three examined prepositions [Wermter 89] showed that the backpropagation networks learned almost all prepositional relationships in the training set and most of the relationships in the testing set. For instance, the network for the preposition "in" got 93% of the 248 training examples correct and 83% of 30 unknown testing examples.

### 3.2 The Bottom Level: Representing Syntactic Constraints

The first form of syntactic knowledge considered for noun phrase disambiguation is the locality constraint. The **Locality constraint** models the heuristic that a prepositional phrase in a noun phrase is more likely to attach to a close preceding noun than to a distant preceding noun. For instance in a noun phrase like "techniques for measurements in discharges" the prepositional phrase "in discharges" tends to attach to "measurements", although "in discharges" could attach to "techniques" as well. The locality constraint can be interpreted as a generalization of Right Association [Kimball 73]. While Right Association for a noun phrase states that a prepositional phrase attaches to the directly preceding noun, the locality constraint claims that there is only a strong tendency for a local attachment to directly preceding nouns. This tendency decreases with the distance between noun and prepositional phrase.

The second form of syntactic knowledge is the **No-crossing constraint**. The no-crossing constraint states that the prepositional phrase attachment in a noun phrase does not show crossing branches (see e.g. [Tait 83]). The following constructed example illustrates a violated no-crossing constraint. Although "influence on electrons" and "temperatures in Fahrenheit" are plausible prepositional relationships, this structural interpretation is considered wrong due to the crossing attachment.

Influence of temperatures on electrons in Fahrenheit.

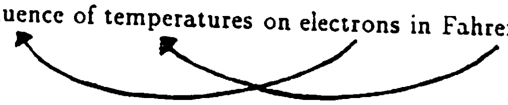

### 3.3 The Top Level: Integrating Syntactic and Semantic Constraints in a Relaxation Network

The semantic constraints for the prepositional relationships and the syntactic constraints for no-crossing and locality are integrated in a relaxation network to allow parallel interactions between these different constraints. In the past, relaxation networks have been shown to be successful for integrating different constraints in a variety of natural language tasks like sentence processing [Waltz and Pollack 85], word sense disambiguation [Bookman 87], attachment decisions [Lehnert 89] and lexical access [Cottrell 88]. These approaches depend on the initialization of the input nodes with suitable values but this decision is not based on a memory model. In our new approach we demonstrate that (1) trained backpropagation networks supply a more powerful underlying model for the input of a relaxation network and (2) relaxation networks are extremely useful for integrating different constraints for structural noun phrase disambiguation.

First we will describe the interface between our two levels, then we will outline the overall architecture of the relaxation network at the top level. This description is illustrated with an example of a noun phrase with three prepositions:

The influence of the radiation intensity on discharges in the Van-Allen-belt.

#### 3.3.1 The Interface between the Top and Bottom Levels

The interface between the two levels is represented with three types of nodes: semantic nodes, locality nodes, and no-crossing nodes (see figure 3). In our example there are six **Semantic nodes** representing the semantic constraints for the six possible prepositional relationships: influence of intensity, influence on discharges, influence in Van-Allen-belt, intensity on discharges, intensity in Van-Allen-belt, discharges in Van-Allen-belt.

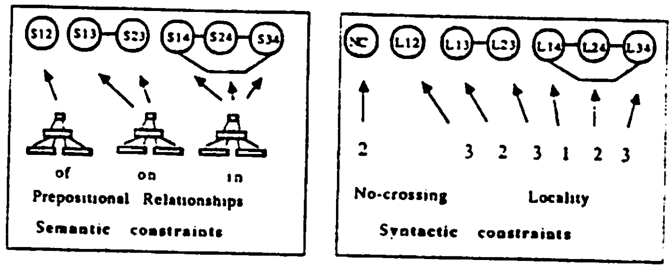


Noun phrase: Influence of intensity on discharges in Van-Allen-Belt

12: influence of intensity
23: intensity on discharges
13: influence on discharges
34: discharges in Van-Allen-Belt
24: intensity in Van-Allen-Belt
14: influence in Van-Allen-Belt

Figure 3: The Interface between Top Level and Bottom Level: Input Nodes for Semantic and Syntactic Constraints

The input potential for the six semantic nodes in the relaxation network is based on the output units of the backpropagation networks described in section 3.1. The semantic nodes representing "influence of intensity", "influence on discharges", "influence in Van-Allen-belt", "intensity in Van-Allen-belt", and "discharges in Van-Allen-belt" get high input potential, because these relationships are plausible. The semantic node for "intensity on discharges" gets a low input potential, because that relationship is implausible.

In addition to the semantic nodes there are seven syntactic nodes representing the syntactic constraints for locality and crossing dependencies. The potential of the six **Locality nodes** reflects the distance between the nouns of a prepositional relationship: the closer the nouns of a prepositional relationship in the noun phrase, the higher the potential of the node. For instance, "influence of intensity" gets a higher value than "influence in Van-Allen-belt" because the nouns in the first prepositional relationship are closer. The one **No-crossing node** prevents crossing attachments, so that in noun phrases with three prepositions the third noun cannot attach to the first noun while the fourth noun attaches to the second. The connections of all nodes are described in the next section.

#### 3.3.2 The Top Level: Architecture of the Relaxation Network

The relaxation network (see figure 4) consists of nodes connected via inhibitory and excitatory connections and can be generated for noun phrases with different lengths. For noun phrases with three prepositions there are 13 input nodes and six output nodes. The input nodes for the semantic constraints and locality constraints are connnected via inhibitory connections if the two prepositional relationships have the same noun in the second position of the prepositional relationship and a different noun in the first position. For example "influence on discharges" and "intensity on discharges" are connected via inhibitory connections, because "influence" competes with "intensity" for "discharges".

The output nodes represent the six possible structural interpretations of the noun phrase. Therefore the output nodes will be referred to as **Structure nodes**. One structure node can be described as a triple of number pairs. Each number stands for the position of a noun in a noun phrase, for instance the triple "1-2,2-3,3-4" is the representation for "influences of intensity", "intensity on discharges" and "discharges in Van-Allen-belt". All structure nodes are in competition and connected via inhibitory connections.

The semantic nodes and the locality nodes are connected with the structure nodes via excitatory connections if the prepositional relationship of the input node occurs in the structure node. The no-crossing node is inhibitorily connected to the structure node "1-2,1-3,2-4" which represents crossing dependencies.

#### 3.3.3 Processing in the Relaxation Network

The nodes in the relaxation network are initialized with a potential between 0 and 10. The semantic nodes receive input based on the output of the backpropagation networks. They obtain a high start potential of 10 for a plausible prepositional relationship and a low start potential of 2 for an implausible prepositional relationship. The initialization values of the locality nodes depend on the distance between nouns in a noun phrase. For instance, if the attachment is over 1 preposition we initialize with 3, for attachment over 2 prepositions with 2, and for attachment over 3 prepostions with 1. This ensures that local attachment gets more reinforcement than distant attachment. The rest of the nodes, the no-crossing node and the structure nodes, are initialized with low values of 2.

Once the relaxation algorithm [Feldman and Ballard 82] is started, nodes update their potential. Incoming excitatory connections increase the potential of a node, incoming inhibitory connections decrease the potential. One cycle consists of updating every node once. Although our implementation of this process is sequential, the actions within one cycle could be processed in parallel. After about 30 cycles the network converges to a stable state in which the potentials do not change any more. The structure node with the highest potential represents the preferred structural interpretation of the noun phrase.

In our example "The influence of the radiation intensity on discharges in the Van-Allen-belt" the following structure node had the highest potential of 8.9 at the end of the relaxation (the other structure nodes had values around 0.9):

Influence of intensity
Influence on discharges
Discharges in Van-Allen-belt

The network integrated the syntactic and semantic constraints: the semantic constraint "intensity on discharges" is implausible and therefore the semantic constraint "influence on discharges" is found as the preferred attachment for "discharges", although the syntactic locality constraint prefers the local attachment "intensity on discharges" compared to "infuence on discharges". This example shows how semantic constraints can overrule locality constraints.

Looking at the noun "Van-Allen-belt" we notice the syntactic influence. "Van-Allen-belt" could attach to all three preceding nouns, because all these prepositional relationships are plausible. At the same time the locality constraint imposes a preference for a local attachment, so that "discharges in Van-Allen-belt" is preferred to "influence in Van-Allen-belt" and "intensity in Van-Allen-belt".

## 4 Discussion

We use two different mechanisms at two levels for the task of structural noun phrase disambiguation. At the domain-dependent bottom level we use distributed representations and backpropagation networks for each preposition to learn the semantic relationships. At the domain-independent top level we use localist representations and a relaxation network to integrate syntactic and semantic constraints. Although work on related relaxation networks has to rely on some initial setting of the start activation (e.g. [Waltz and Pollack 85], [Bookman 87], [Lehnert 87], [Cottrell 88]), our model bases its initialization on learned memory. While other work on PP-attachment has mostly concentrated on the attachment of single prepositional phrases in sentences ([Wilks et al. 85], [Schubert 86], [Dahlgren and McDowell 86], [McClelland and Kawamoto 86], [St. John and McClelland 88]) we have concentrated on the attachment of multiple prepositional phrases in noun phrases.

Our approach demonstrates progress over related connectionist work [Cosic and Munro 88] by using distributed representations for nouns, by integrating semantic and syntactic constraints and by allowing for noun phrases with arbitrary length. We must also point out that our underlying memory model of prepositional relationships can be used as part of a full sentence ana-

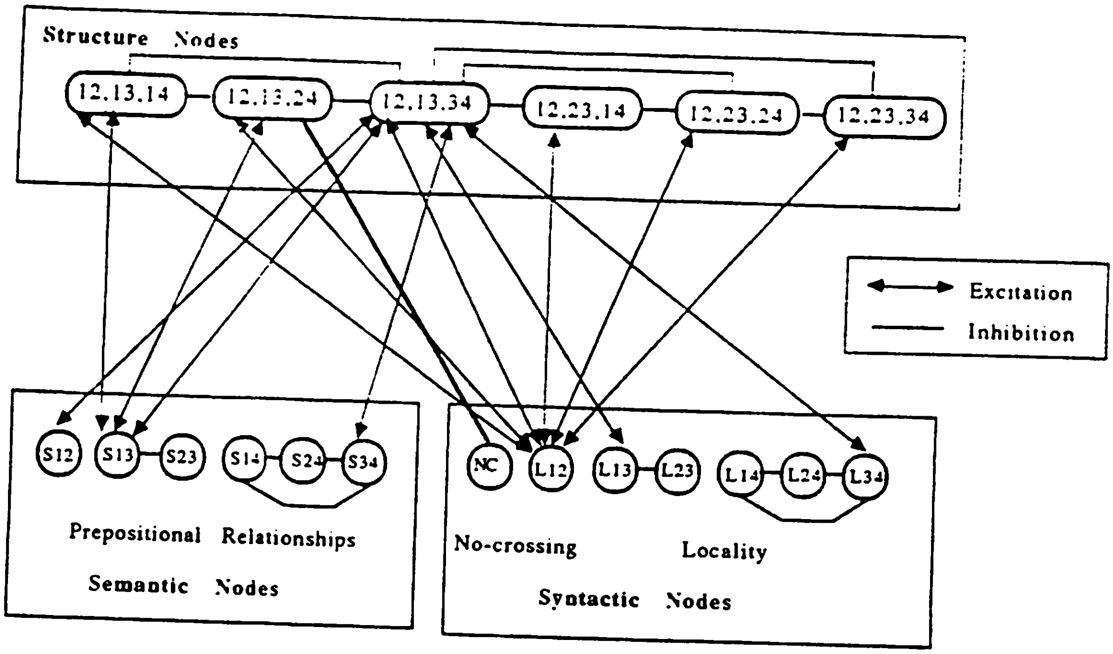


Noun phrase: Influence of intensity on discharges in Van-Allen-Belt

12: influence of intensity
23: intensity on discharges
13: influence on discharges

34: discharges in Van-Allen-Belt
24: intensity in Van-Allen-Belt
14: influence in Van-Allen-Belt

Connections: Only the excitatory connections for the semantic node S13, for the syntactic node L12 and for the structure node 12.13.34 are shown completely

Only the inhibitory connections for structure node 12.13.34 are shown completely

Figure 4: Top Level: Relaxation Network for the Integration of Semantic and Syntactic Constraints (only some connections shown)

lyzer as well. For example, the sentence analyzer CIRCUS [Lehnert 89] can combine our semantic memory model with predictive knowledge during sentence processing.

## 5 Conclusions

We have described an approach for learning and integrating semantic and syntactic constraints. Backpropagation networks and distributed representations are used to learn the plausibility of semantic relationships and to generalize the learned regularities to semantic constraints. Relaxation networks and localist representations are used to integrate these semantic constraints with syntactic constraints. We have demonstrated that a connectionist model supplies a powerful memory model for the learning and integration of constraints for structural noun phrase disambiguation. Since the problem of learning and integrating constraints occurs in many other language tasks like word sense disambiguation or compound noun interpretation, our memory model is of importance for many Natural Language Processing problems.


## Acknowledgements

I would like to thank Wendy Lehnert, Andy Barto, Steven Bradtke, David Lewis and the NLP group for discussions and comments on earlier drafts of this paper.


## References


1. Bookman L.A. 1987. A microfeature based scheme for modelling semantics. *Proceedings of the IJCAI*
2. Cottrell G.W. 1988. A Model of Lexical Access of ambiguous words. In: Small, S.L., Cottrell,G.W., Tanenhaus,M.K. *Lexical Ambiguity Resolution*
3. Cosic C., Munro P. 1988. Learning to represent and understand locative prepositional phrases. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society.*
4. Dahlgren K., McDowell J. 1986. Using Commonsense Knowledge to disambiguate Prepositional Phrase Modifiers. *Proceedings of the AAAI.*
5. Feldman J.A., Ballard D.H. 1982. Connectionist Models and their properties. *Cognitive Science 6*

6. Herskovits A. 1986. *Language and spatial cognition.* Cambridge University Press.

7. Kimball J. 1973. Seven principles of surface structure parsing in natural language. *Cognition 2.*

8. Lehnert W.G. 1987. Learning to integrate syntax and semantics. *Proceedings of the Fourth International Workshop on Machine Learning.*

9. Lehnert W.G. 1989. Symbolic/Subsymbolic Sentence Analysis: Exploiting the Best of Two Worlds. In: Barnden,J., Pollack,J.: *Advances in Connectionist and Neural Computation Theory* Vol 1, Ablex Publishers, in press (also available as COINS Technical Report 88-99, Department of Computer and Information Science, University of Massachusetts, Amherst, MA)

10. McClelland J.L., Kawamoto A.H. 1986. Mechanisms of Sentence Processing: Assigning Roles to Constituents. In: McClelland,J.L., Rumelhart,D.E. (Ed.) *Parallel Distributed Processing Vol 2.* MIT Press.

11. Rumelhart D.E., Hinton G.E., Williams R.J. 1886. Learning Internal Representations by Error Propagation. In: McClelland,J.L., Rumelhart,D.E. (Ed.) *Parallel Distributed Processing Vol 1.* MIT Press.

12. Schubert L.K. 1986. Are there preference tradeoffs in Attachment decisions? *Proceedings of AAAI.*

13. Sparck-Jones K., Van Rijsbergen C.J. 1976. Progress in Documentation. *Journal of Documentation, vol 32, no 1*

14. St. John M.F., McClelland J.L. 1988. Applying contextual constraints in sentence comprehension. *Proc. of the Tenth Annual Conference of the Cognitive Science Society.*

15. Tait J.I. 1983. Semantic parsing and syntactic constraints (Mark IV). In: Sparck Jones K., Wilks Y.: *Automated natural language processing.* Ellis Horwood Limited.

16. Waltz D.L., Pollack J.B. 1985. Massively Parallel Parsing: A strongly interactive model of natural language interpretation. *Cognitive Science 9*

17. Wermter S., Lehnert W.G. 1989. Noun phrase analysis with connectionist networks. In: Sharkey N., Reilly R.: *Connectionist Approaches to Language Processing.* (in press).

18. Wilks Y., Huang X.,. Fass D. 1985. Syntax, Preference and Right Attachment. *Proceedings of the IJCAI.*

# Learning Semantic Relationships in Compound Nouns with Connectionist Networks

Stefan Wermter
Department of Computer and Information Science
University of Massachusetts

## Abstract

This paper describes a new approach for understanding compound nouns. Since several approaches have demonstrated the difficulties in finding detailed and suitable semantic relationships within compound nouns, we use only a few basic semantic relationships and provide the system with the additional ability to learn the details of these basic semantic relationships from training examples. Our system is based on a backpropagation architecture and has been trained to understand compound nouns from a scientific technical domain. The test results demonstrated that a connectionist network is able to learn semantic relationships within compound nouns.

## Introduction

Understanding compound nouns plays an important role in understanding natural language. In the past, different approaches for understanding compound nouns have been investigated in artificial intelligence, linguistics, and cognitive science ((Marcus 80) (Finin 80) (McDonald 82) (Lehnert 86) (Arens 87) (Dahl 87)). Most approaches relied on a representation of the words in compound nouns as frames or semantic features and contained fixed control structures which determined the semantic relationships between the words. For example, Finin (Finin 80) used frames to predict the semantic relationships between words and a hierarchy of rules to identify the best relationship. McDonald's system (McDonald 82) is based on Fahlman's parallel semantic network (Fahlman 79) and used marker passing to find the semantic relationships between word concepts.

These approaches try to understand compound nouns by coding as much knowledge as possible about the words, semantic relationships, and control structures. In this paper we investigate a different approach for understanding compound nouns consisting of two words. We use only a few basic semantic relationships and provide the system with the ability to learn the details of the basic semantic relationships from training examples. Instead of encoding knowledge structures and control structures for understanding compound nouns, basic semantic relationships in compound nouns are learned using a connectionist architecture.

## The Domain and the Basic Semantic Relationships

Compound nouns are frequently used in almost every domain. Our domain is the NPL[1] corpus (Sparck-Jones 76) which contains abstracts and queries from the physical sciences. From this corpus

[1] National Physics Laboratory

we randomly chose 108 compound nouns consisting of two words, e.g. "heat effect". Each word is represented as a binary vector of 16 semantic features, which were extracted by using the NASA thesaurus (NASA 85). For a more detailed description of the process of feature extraction see (Wermter and Lehnert 89). Figure 1 illustrates the semantic features for the compound nouns.

| Semantic Features | Examples |
|---|---|
| MEASURING-EVENT | Observation, Investigation, Research |
| CHANGING-EVENT | Amplification, Acceleration, Loss |
| SCIENTIFIC-FIELD | Mechanics, Ferromagnetics |
| PROPERTY | Intensity, Viscosity, Temperature |
| MECHANISM | Experiment, Technique, Theorem |
| ELECTRIC-OBJECT | Transistor, Resistor, Amplifier |
| PHYSICAL-OBJECT | Earth, Crystal, Vehicle, Room |
| RELATION | Cause, Dependence, Interaction |
| ORGANIZATION-FORM | Layer, Level, Stratification, F-Region |
| GAS | Air, Oxygen, Atmosphere, Nitrogen |
| SPATIAL-LOCATION | Antarctic, Earth, Range, Region, Source |
| TIME | June, Day, Time, History |
| ENERGY | Radiation, Ray, Light, Sound, Current |
| MATERIAL | Aluminium, Water, Carbon, Vapour |
| ABSTRACT-REPRESENTATION | Note, Data, Equation, Term, Parameter |
| EMPTY | Cavity, Vacua |

Figure 1: Semantic Features of the Nouns and Examples

To represent basic semantic associations between words we use 7 basic semantic relationships. We specify a **Basic Semantic Relationship** as a preposition paraphrase (see figure 2). For example, a "room experiment" has the basic semantic relationship IN-P since the experiment is "in" a room, and an "excitation mechanism" has the basic semantic relationship FOR-P since it is a mechanism "for" excitation. Each compound noun can have different basic semantic relationships; for instance, a "feedback circuit" is a "circuit FOR-P feedback" or a "circuit WITH-P feedback". Each basic semantic relationship can have several meanings; for instance, the IN-P is different for "storage IN-P computer" and "disturbance IN-P atmosphere".

| Basic Semantic Relationships | Examples for the Basic Semantic Relationships |
|---|---|
| BY-P | Impurity Conduction |
| FOR-P | Excitation Mechanism |
| FROM-P | Space Vehicle |
| IN-P | Room Experiment |
| OF-P | Oxygen Emission |
| ON-P | Skin Effect |
| WITH-P | Amplifier Circuit |

Figure 2: The Basic Semantic Relationships

We consider the basic semantic relationships as a first step to differentiate semantic relationships according to their main properties. This general concept of classifying semantic relationships according to preposition paraphrases has been found useful in several studies on compound nouns (e.g., (Lee 60) (Levi 78) (Finin 80)), since preposition paraphrases contain general relationships; e.g., FROM-P expresses a source, FOR-P expresses a purpose, and IN-P expresses inclusion. Our goal here is to specify basic semantic relationships as preposition paraphrases and to build a system which learns the underlying semantic relationships from training examples.

## The Architecture

The architecture for learning semantic relationships is a backpropagation network with three layers (see figure 3). The bottom layer consists of 32 binary input units for the semantic features of the two words in the compound noun. The hidden layer is a 7 x 12 array of hidden units, 12 hidden units for each of the 7 basic semantic relationships. The top layer consists of 7 real-valued output units, one for each of the 7 basic semantic relationships.

Each output unit is connected only to all hidden units of the same basic semantic relationship. All hidden units are connected to all input units. This modular organization has two advantages: (1) training and testing for each basic semantic relationship can be done independently, and (2) adding, deleting and modifying a basic semantic relationship does not require retraining the whole network.

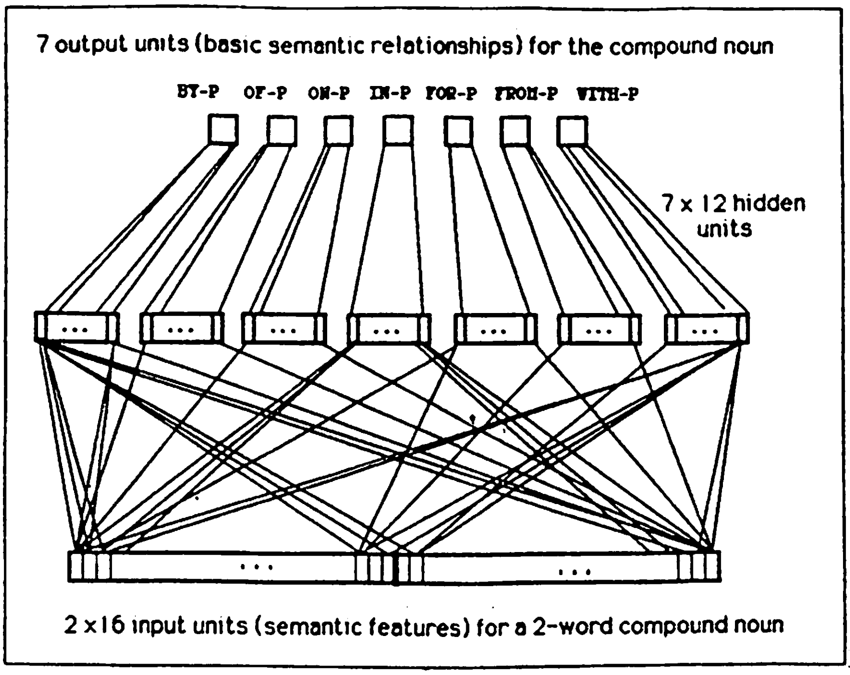


Figure 3: The Structure of the Backpropagation Network

## Training the Network

First, 108 compound nouns consisting of two words were randomly selected from the NPL corpus. Each compound noun was represented with 32 binary features, 16 for each word. The 108 compound nouns were divided into 88 compound nouns for a training set and 20 compound nouns for a test set. Because of the modular architecture the network can be trained in separate modules for the different basic semantic relationships. For each of the 7 basic semantic relationships the feature representations of the 88 compound nouns were presented as the input together with a desired binary plausibility value as the output. The plausibility value indicates if the basic semantic relationship between the two words is plausible (value 1) or not plausible (value 0). The following example shows two of the 88 training examples for the basic semantic relationship IN-P: "Plasma layer" in the sense of "layer IN-P plasma" is plausible, while "sunspot number" in the sense of "number IN-P sunspot" is not plausible.

```
PLASMA LAYER      ->     LAYER IN-P PLASMA        1
SUNSPOT NUMBER    ->     NUMBER IN-P SUNSPOT      0
```

For each of the 7 basic semantic relationships the semantic features and plausibility values of the 88 compound nouns were presented for 800 cycles (that is 70400 training examples). The backpropagation algorithm (Rumelhart et. al. 86) was used to learn the plausibility of each basic semantic relationship[2]. To be independent of the random start initialization of the network, three different runs (each with the 70400 training examples) were conducted for each of the 7 basic semantic relationships. Within this learning phase the average of the total sum squared error for all training examples over all 21 runs decreased from 23.2 at the start of the training to 1.4 at the end of the training.

## Evaluation of the Test Results

After training, the network was tested on the training set of 88 compound nouns and the test set of 20 compound nouns. The semantic feature representation of the compound nouns in the test set had not been part of the training set. The network was tested by presenting the feature representation of a compound noun, and the system computed the plausibility value for each basic semantic relationship. A basic semantic relationship is considered correct, if the computed plausibility value deviates less than 0.49 from the desired value 1 for a plausible basic semantic relationship and from the desired value 0 for an implausible basic semantic relationship.

[2]The learning rate $\eta$ was set to 0.01, the weight change momentum $\alpha$ was 0.9 for all experiments.

| Basic Semantic Relationships | Correct in the Training Set | Correct in the Test Set |
|---|---|---|
| BY-P | 94% | 83% |
| FOR-P | 97% | 73% |
| FROM-P | 94% | 82% |
| IN-P | 96% | 73% |
| OF-P | 93% | 77% |
| ON-P | 98% | 95% |
| WITH-P | 98% | 88% |

Figure 4: Basic Semantic Relationships in Training Set and Test Set

Figure 4 illustrates the overall system performance on the training set and on the test set for each basic semantic relationship. The average percentage of correctly learned training examples for the three different learning runs is between 93% and 98%, the percentage of correctly generalized test examples is between 73% and 95%.

Figure 5 shows a more detailed interpretation of representative examples from the test set of new compound nouns. Each compound noun is shown with the computed plausibility values for each basic semantic relationship[3]. We say that a basic semantic relationship for a compound noun exists if the computed plausibility value is greater than or equal to 0.5.

| Compound Nouns | BY-P | FOR-P | FROM-P | IN-P | OF-P | ON-P | WITH-P |
|---|---|---|---|---|---|---|---|
| Heat Exchange | 0.3 | 0.0 | 0.0 | 0.3 | 0.9 | 0.0 | 0.0 |
| Transistor Life | 0.0 | 0.4 | 0.0 | 0.1 | 1.0 | 0.0 | 0.1 |
| Writing Method | 0.0 | 0.8 | 0.1 | 0.0 | 1.0 | 0.0 | 0.0 |
| Wing Motion | 0.0 | 0.1 | 0.3 | 1.0 | 0.5 | 0.0 | 0.0 |
| Waveform Solution | 0.1 | 0.4 | 0.0 | 0.1 | 0.4 | 0.1 | 0.0 |
| Earth Satellite | 0.0 | 0.0 | 0.9 | 0.9 | 0.0 | 0.0 | 0.7 |
| Transport Theory | 0.3 | 0.1 | 0.0 | 0.0 | 0.9 | 0.6 | 0.1 |
| Water Vapour | 0.0 | 0.0 | 0.6 | 1.0 | 0.4 | 0.0 | 0.6 |
| Wave Propagation | 0.1 | 0.0 | 0.2 | 0.7 | 0.7 | 0.1 | 0.0 |
| Microwave Emission | 0.7 | 0.1 | 0.0 | 0.0 | 1.0 | 0.0 | 0.0 |

Figure 5: Examples for the Interpretation of Compound Nouns (see text for explanation)

In the first two examples in figure 5 a single basic semantic relationship exists between the two words in the compound noun: "heat exchange" is interpreted as "exchange OF-P heat", and "transistor life" as "life OF-P transistor" (only these basic semantic relationships have a plausibility value greater or equal 0.5).

[3]Again, as in figure 4, the plausibility values shown are the averages over the three different runs for each basic semantic relationships.

Although only one basic semantic relationship exists in the first two examples, most test examples have more than one existing basic semantic relationship. For instance, "writing method" has the existing relationships "method FOR-P writing" and "method OF-P writing". The other basic semantic relationships for "writing method", like "method BY-P writing" and "method FROM-P writing", do not exist. Another example of multiple basic semantic relationships is "wing motion" (as in airplanes) which is interpreted as "motion OF-P wing" and "motion IN-P wing". This example illustrates ambiguous interpretations and context is needed to determine if the wing is the object which is moving (motion OF-P wing) or the location of a motion (motion IN-P wing).

The plausibilty values in Figure 5 indicate unsure interpretations as well. For instance, the plausibility values of the compound noun "waveform solution" are lower than 0.5 for all basic semantic relationships. The network can not find a basic semantic relationship because similar relationships had not been in the training set. The results show examples with some incorrect basic semantic relationships as well. For instance "water vapour" is interpreted with 3 existing relationships: "vapour FROM-P water", "vapour WITH-P water", and "vapour IN-P water". While the first two relationships FROM-P and WITH-P are plausible, the third is not plausible.

Although our corpus is still fairly small our test results demonstrate the extent to which the learned basic semantic relationships generalize for new compound nouns. The basic semantic relationships in our network generalize well for compound nouns whose first and second noun are characterized with subsets of the following semantic features: Noun1: ENERGY PROPERTY ORGANIZATION-FORM and Noun2: CHANGING-EVENT PROPERTY MECHANISM. Examples for this class of compound nouns are "heat exchange" and "wave propagation". Another class of compound nouns with good generalizations are subsets of the following features: Noun1: ELECTRIC-OBJECT PHYSICAL-OBJECT and Noun2: TIME PROPERTY, like in "transistor life".

Besides these classes of compound nouns with good generalizations, compound nouns with subsets of the following features do not generalize well: Noun1: PHYSICAL-OBJECT SPATIAL-LOCATION MATERIAL and Noun2: PHYSICAL-OBJECT GAS MATERIAL. Examples with subsets of these feature combinations are "earth satellite" and "water vapour". The reason for the decrease in the generalization performance for this last class is the restricted use of only 16 semantic features. To generalize relationships between two physical objects more features are needed. For instance, a network with a SIZE feature could generalize the WITH-P relationships between physical objects so that "earth satellite" could not be interpreted as "satellite WITH-P earth" since the earth has a bigger size than a satellite. The identification of these incorrectly generalized basic relationships is important for deciding which semantic features and basic semantic relationships might be modified. We make no claim for a "right" classification of semantic features and basic semantic relationships for our domain but we claim that the adaptive process of identifying better suitable semantic features and semantic relationships is supported by the learning ability and the modular architecture.

## Related work

Comparing the performance of our system with existing systems for compound noun analysis is somewhat difficult, because the techniques, the level of the semantic relationships, and the domains are fundamentally different. McDonald reports about 54% to 64% correct interpretations for his compound noun system (McDonald 82) using detailed semantic relationships and fixed control strategies. The performance of Finin's system is similar to McDonald's system. Our system determines plausible basic semantic relationships for unknown compound nouns. Although our basic semantic relationships are not as detailed as McDonald's or Finin's, our basic semantic relationships are automatically acquired. As far as we know there is currently no system which has the ability to *learn* the semantic relationships between compound nouns.

Our system has the advantage of learning knowledge for the semantic relationships, while this knowledge is difficult to acquire in other compound noun systems (e.g. (Finin 80) (McDonald 82) (Arens 87) (Gay 88)). The knowledge about semantic relationships is represented uniformly in modular networks. On the other hand, these systems allow compound nouns with more than two words while we need additional mechanisms to understand longer compound nouns. Currently, we are investigating the use of recursive autoassociative network architectures ((Pollack 88), (St John 88)) and relaxation networks (Wermter 89) to understand compound nouns of arbitrary length.

## Conclusions

One way to approach compound noun analysis is the use of extensive knowledge engineering, as demonstrated in several computational models. Because of the difficulties of identifying the semantic relationships and control structures, we presented a new approach for understanding compound nouns. Using a modular connectionist architecture we showed that basic semantic relationships within compound nouns can be learned. The general concepts of basic semantic relationships, learning, and modular network architectures demonstrate how uniform memory models can be built for natural language understanding.

## Acknowledgements

I would like to thank Wendy Lehnert, Steve Bradtke, Claire Cardie, and Ellen Riloff for discussions and comments on earlier drafts of this paper. This research was supported by the Advanced Research Projects Agency of the Department of Defense, monitored by the Office of Naval Research under contract #N00014-87-K-0238, the Office of Naval Research, under a University Research Initiative Grant, Contract #N00014-86-K-0764 and NSF Presidential Young Investigators Award NSFIST-8351863.

## References

1. Arens Y., Granacki J.J., Parker, A.C. 1987. Phrasal Analysis of Long Noun Sequences. *Proceedings of the ACL 87.*

2. Fahlman S.E. 1979. *NETL: A system for representing and using real-world knowledge.*

3. Finin T.W. 1980. *The semantic Interpretation of Compound Nominals.* PhD Thesis. University of Illinois at Urbana-Champaign.

4. Dahl A.D., Palmer M.S., Passonneau R.J. 1987. Nominalizations in PUNDIT. *Procedings of the ACL 1987.*

5. Gay L.S. 1988. *Interpreting Nominal Compounds for Information Retrieval.* Internal Report, GE Research and Development Center.

6. Lee R.B. 1960. The grammar of English Nominalizations. *International Journal of American Linguistics, Vol. 26, 3.*

7. Lehnert W.G. 1986. The analysis of compound nouns. *VERSUS, December, 1986.*

8. Levi J.N. 1978. The syntax and semantics of complex nominals.

9. Marcus M.P. 1980. *A theory of syntactic recognition for natural language.*

10. McDonald D.B. 1982. *Understanding Noun Compounds.* PhD thesis, Carnegie Mellon University.

11. NASA Thesaurus 1985.

12. Pollack J. 1988. Recursive Auto-associative Memory: Devising Compositional Distributed Representations. *Proc. of the Tenth Annual Conference of the Cognitive Science Society.*

13. Rumelhart D.E., Hinton G.E., Williams R.J. 1886. Learning Internal Representations by Error Propagation. In: McClelland J.L., Rumelhart D.E. (Ed.) *Parallel Distributed Processing Vol 1.* MIT Press.

14. Sparck-Jones K., VanRijsbergen C.J. 1976. Information Retrieval Test Collections. *Journal of Documentation, vol 32, no 1.*

15. St.John M.F., McClelland J.L. 1988. Applying contextual constraints in sentence comprehension. *Proc. of the Tenth Annual Conference of the Cognitive Science Society.*

16. Wermter S., Lehnert W.G. 1989. Noun phrase analysis with connectionist networks. In: Sharkey N., Reilly R.: *Connectionist Approaches to Language Processing.* (in press).

17. Wermter S. 1989. Integration of semantic and syntactic constraints for structural noun phrase disambiguation. *Proc. of the Eleventh International Joint Conference on Artificial Intelligence 89.*

# A Hybrid Symbolic/Connectionist Model for Noun Phrase Understanding

Stefan Wermter, Wendy G. Lehnert
Department of Computer and Information Science
University of Massachusetts at Amherst
Amherst, MA 01003
USA

## Abstract

This paper describes a hybrid model which integrates symbolic and connectionist techniques for the analysis of noun phrases. Our model consists of three levels: (1) a distributed connectionist level, (2) a localist connectionist level, and (3) a symbolic level. While most current systems in natural language processing use techniques from only one of these three levels, our model takes advantage of the virtues of all three processing paradigms. The distributed connectionist level provides a learned semantic memory model. The localist connectionist level integrates semantic and syntactic constraints. The symbolic level is responsible for restricted syntactic analysis and concept extraction. We conclude that a hybrid model is potentially stronger than models that rely on only one processing paradigm.

# 1. Introduction

In recent years there has been a growing interest in using connectionist techniques for natural language processing. While traditionally the analysis, representation, and generation of natural language were exclusively dominated by symbolic approaches, lately connectionist techniques have received increased attention because of their attractive properties including noise resistance, learning behavior, neural plausibility, associative retrieval, and knowledge integration.

There have been at least two main directions of work in connectionist artificial intelligence: implementation-oriented and task-oriented. Implementation-oriented connectionism tries to show how symbolic representations and computations can be implemented with connectionist techniques. Connectionist systems have been developed to implement semantic networks (Shastri 1988), rule-based systems (Touretzky & Hinton 1988) (Shastri & Ajjanagadde 1989), representation languages like KL-ONE (Derthick 1988), hierarchies and tree-like structures (Hinton 1988) (Pollack 1988). Other connectionist systems show how symbolic computations can be implemented, e.g., variable binding (Touretzky & Hinton 1985), sequential processing (Jordan 1986) (Elman 1988), and recursion (Pollack 1989). This implementation-oriented research demonstrates that connectionist models can, at least to a certain extent, implement symbolic structures and computations.

Task-oriented connectionism tries to show how specific tasks can be modeled with connectionist techniques. Numerous tasks in natural language processing have been attacked in recent years, e.g., parsing (Fanty 1985) (Hanson & Kegl 1987) (Howells 1988) (Kwasny 1988), word sense disambiguation (Cottrell & Small 1983) (Bookman 1987), anaphor resolution (Allen 1987), compound noun understanding (Wermter 1989b), sentence generation (Gasser 1988), script and concept understanding (Dolan & Dyer 1988) (Miikkulainen & Dyer 1989), language acquisition (Rumelhart & McClelland 1986), and role assignment (McClelland & Kawamoto 1986) (St. John & McClelland 1988). These approaches demonstrate that connectionist models are useful for certain restricted tasks in natural language processing.

Implementation-oriented and task-oriented connectionism both demonstrate several advantages and disadvantages of symbolic and connectionist processing techniques. Although purely connectionist systems (Waltz & Pollack 1985) (Sejnowski & Rosenberg 1986) (Hanson & Kegl 1987) and purely symbolic systems (Charniak 1983) (Dyer 1983) (Riesbeck & Martin 1986) (Grosz et al. 1987) (Hirst 1987) have both shown impressive results, it has become obvious that connectionist techniques and symbolic techniques exhibit complementary strengths (Touretzky 1988) (Lehnert 1988) (Dyer 1988) (Hendler 1989). While symbolic processing has advantages in representing schemata, recursive structures, variable binding, inheritance hierarchies, and sequential control, connectionist processing has advantages in associative retrieval, noise resistance, knowledge integration, generalization, and learning. Because of these mutually complementary properties, hybrid symbolic/connectionist systems promise to be more powerful than systems operating within only one paradigm.

In this paper we present a hybrid model for understanding noun phrases. This model combines localist and distributed connectionist techniques with symbolic techniques. The model consists of three levels: (1) distributed connectionist networks are used to learn semantic relationships between nouns, (2) localist connectionist networks integrate semantic constraints and syntactic constraints, and (3) symbolic techniques provide a restricted syntactic analysis and concept extraction. In section 2 we describe our domain and our noun representation, in section 3 the distributed connectionist level, in section 4 the localist connectionist level, and in section 5 the symbolic level. We show how a hybrid model can be used to understand noun phrases from a scientific technical domain.

## 2. The Domain: Noun Phrases in Scientific and Technical Sublanguages

Noun phrases are the dominant source of information in scientific and technical sublanguages (Hirschman 1986). Because noun phrases are so important, natural language processing systems in these domains need a powerful and flexible model for understanding noun phrases. To investigate such a model we chose noun phrases from the NPL[1] corpus (Sparck Jones & Van Rijsbergen 1976) as our domain. The NPL corpus contains queries and titles of scientific articles from the physical sciences. For example:

Effects of electromagnetic fields on turbulences in gases.

Note on the cause of ionization in the f-region.

[1] National Physics Laboratory.

Radio emission by plasma oscillations in nonuniform plasmas.

Calculation of fields on plasma ions by collective coordinates.

An iterative analogue computer for use with resistance network analogues.

Syntactic, semantic, contextual, and world knowledge are all necessary for understanding complex noun phrases containing multiple prepositional phrases. In the past, several techniques have been developed to describe the problem of attaching prepositional phrases to the correct constituents (Prepositional Phrase Attachment), for instance, (Kimball 1973) (Frazier & Fodor 1978) (Ford et al. 1982) (Crain & Steedman 1985) (Wilks et. al 1985) (Dahlgren & McDowell 1986) (McClelland & Kawamoto 1986) (Schubert 1986) (Hirst 1987) (Lehnert 1987). Our hybrid approach is different from these approaches because we integrate distributed connectionist networks, localist connectionist networks, and symbolic techniques for understanding noun phrases.

Now we describe the representation of nouns in our domain of the physical sciences. We represent a noun as a binary vector of 16 features. This feature representation was developed as follows. First, we used thesaurus knowledge (NASA 1985) (EJC 67) for classifying the nouns occuring in the noun phrases. We categorized each noun according to the most general term in the hierarchy that describes the noun. This step abstracted specific nouns like "carbon resistor", "noise fluctuation", and "transistor" to more general terms like "resistor", "variation", and "semiconductor device". Then, we grouped these most general thesaurus terms into 16 classes which form the basis of our feature representation. These 16 features describe the basic meaning of a noun in our domain.

For example, the term "carbon resistor" is dominated by "resistor" at the most general level and "transistor" is dominated by "semiconductor device". "Resistor" and "semiconductor device" belong to the class (and therefore have the feature) ELECTRIC OBJECT. Each noun can have multiple features. For instance, the noun "acceleration" has the features CHANGING-EVENT and ENERGY. Figure 1 shows all features along with examples taken from the corpus. Having described the domain encoding, we now turn to a description of our three-part model.

| Semantic Features | Examples |
|---|---|
| MEASURING-EVENT | Observation, Investigation, Research |
| CHANGING-EVENT | Amplification, Acceleration, Loss |
| SCIENTIFIC-FIELD | Mechanics, Ferromagnetics |
| PROPERTY | Intensity, Viscosity, Temperature |
| MECHANISM | Experiment, Technique, Theorem |
| ELECTRIC-OBJECT | Transistor, Resistor, Amplifier |
| PHYSICAL-OBJECT | Earth, Crystal, Vehicle, Room |
| RELATION | Cause, Dependence, Interaction |
| ORGANIZATION-FORM | Layer, Level, Stratification, F-Region |
| GAS | Air, Oxygen, Atmosphere, Nitrogen |
| SPATIAL-LOCATION | Antarctic, Earth, Range, Region, Source |
| TIME | June, Day, Time, History |
| ENERGY | Radiation, Ray, Light, Sound, Current |
| MATERIAL | Aluminium, Water, Carbon, Vapor |
| ABSTRACT-REPRESENTATION | Note, Data, Equation, Term, Parameter |
| EMPTY | Cavity, Vacuum |

Figure 1: Semantic features of the nouns and examples

# 3. Learning Semantic Prepositional Relationships in Distributed Connectionist Networks

In this section we describe semantic prepositional relationships and examine how they can be learned. Based on our feature representation, backpropagation networks learn underlying regularities of prepositional relationships from a corpus of training noun phrases. The learned regularities are used for attaching prepositional phrases to appropriate constituents within new noun phrases.

## 3.1 Semantic Prepositional Relationships

Within noun phrases, nouns can be connected with prepositions, as in "Symposium on hydrodynamics in ionosphere". Understanding these noun phrases relies on understanding prepositional relationships. A *prepositional relationship* is the semantic relationship between the features of two nouns which are connected by a preposition. Prepositional relationships can be either plausible or implausible. *Plausible prepositional relationships* are possible relationships, such as "symposium on hydrodynamics". *Implausible prepositional relationships* are relationships which are not reasonable. "Symposium in ionosphere" is implausible because symposiums do not take place in the outer atmosphere.

Knowing about the *plausible* prepositional relationships "symposium on hydrodynamics" and "hydrodynamics in ionosphere" and knowing about the *implausible* prepositional relationship "symposium in ionosphere", we must interpret the noun phrase "symposium on hydrodynamics in ionosphere" so that the prepositional phrase "in ionosphere" attaches to "hydrodynamics", but not to "symposium". Since knowledge about the plausibility of the prepositional relationship between two nouns can help to rule out implausible interpretations of the whole noun phrase, we have trained backpropagation networks to learn the plausibility of prepositional relationships.

## 3.2 Learning Semantic Prepositional Relationships with Backpropagation Networks

We use backpropagation networks (Rumelhart et al. 1986) to learn the plausibility of prepositional relationships within noun phrases. For each preposition there is one backpropagation network that determines the plausibility of the prepositional relationships (see figure 2). One network consists of 3 layers of units. The input layer consists of 32 binary units (values 0 and 1) representing 16 features for each of the two nouns. The single real-valued output unit determines whether the prepositional relationship is plausible (value 1) or implausible (value 0). 12 real-valued hidden units encode the mapping from the input units to the output units from a training set. All levels in the backpropagation network are fully connected. We need one training set of prepositional relationships for each preposition.

First we concentrated on the three prepositions "in", "of", and "on". We randomly extracted 50 noun phrases from our corpus which contained only these three prepositions, for instance:

Note on the cause of ionization in the f-region.

International symposium on fluid mechanics in the ionosphere.

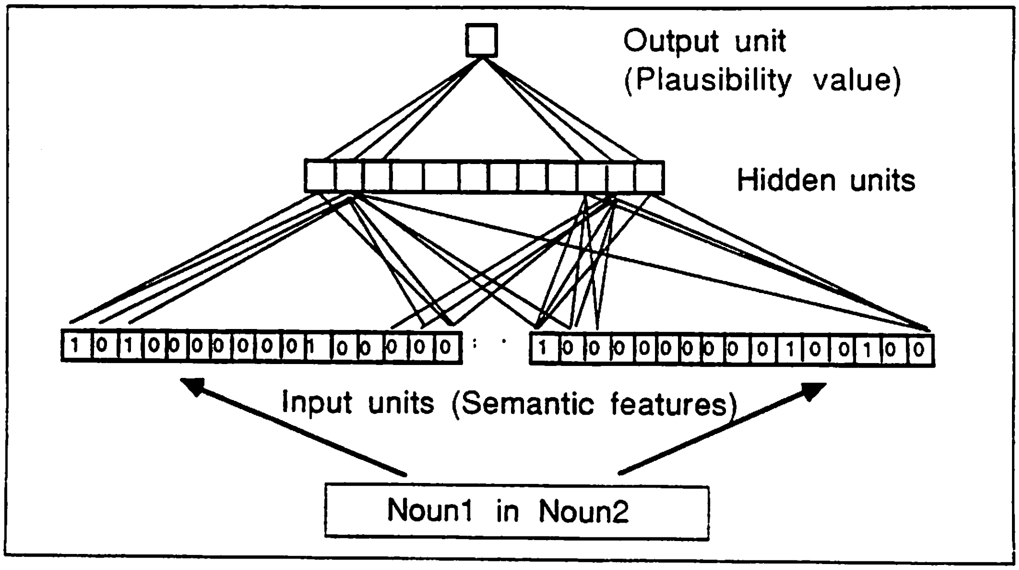


Figure 2: Backpropagation network for the prepositional relationships of "in"

Based on these 50 noun phrases, we built one training set for each preposition. Each training example in the training set consists of two feature vectors for the two nouns together with the binary plausibility value for the prepositional relationship between these nouns. The plausibility value is set to 1 if the prepositional relationship in the training set is plausible and is set to 0 otherwise. From now on and where the context is clear, we will use the term prepositional relationship for both the semantic relationship between the two nouns and the representation of this semantic relationship as a training instance. Each noun in the 50 noun phrases is stored in a lexicon with its name and the associated 16 features. The following examples show two nouns with their features using the same feature order as in figure 1 in section 2.

```
F-region      (0 0 0 0 0 0 1 0 1 1 1 0 0 0 0 0)
Ionization    (0 1 0 1 1 0 0 0 0 0 0 0 0 0 0 0)
```

Now we will describe the training for the prepositional relationships for "in". There were 124 prepositional relationships for the preposition "in" in the 50 noun phrases. Since most of these prepositional relationships in the 50 existing noun phrases are plausible prepositional relationships, most training examples would be plausible prepositional relationships[2]. We added the 124 inverse prepositional relationships to the 124 prepositional relationships so that the training set for "in" consists of 248 prepositional relationships. An inverse prepositional relationship is a prepositional relationship in which the order of the two nouns is changed. Including inverse prepositional relationships in the training set prevents the network from being overloaded with too many plausible relationships since most of the inverse prepositional relationships are implausible. We illustrate the prepositional relationships and the inverse prepositional relationships for the preposition "in" for our example "Note on cause of ionization in f-region" together with their plausibility values:

| Prepositional Relationships | | Inverse Prepositional Relationships | |
|---|---|---|---|
| Note in f-region | 0 | F-region in note | 0 |
| Cause in f-region | 1 | F-region in cause | 0 |
| Ionization in f-region | 1 | F-region in ionization | 0 |

[2]Implausible prepositional relationships like "symposium in ionosphere" in the noun phrase "symposium on hydrodynamics in ionosphere" occur less frequently in existing noun phrases than plausible prepositional relationships.

### 3.2.1 Training results for the prepositional relationships for "in"

Now we show the results for the training set with the 248 prepositional relationships for "in". We conducted three runs training three backpropagation networks with the prepositional relationships for "in". The three different runs show that our training does not depend on a fortuitous initialization of the weights in the network. In each run the backpropagation network was trained for 1600 epochs (396800 prepositional relationships) with the learning rate $\eta = 0.01$ and weight change momentum $\alpha = 0.9$. The weights in the backpropagation network were updated after each complete epoch.

After the training phase was completed, the trained networks were tested with the training set. To interpret the tests we introduce the terms "error tolerance", "error rate", and "total error". The *error tolerance* determines how much the actual outcome of the output unit could deviate from the desired outcome 0 for an implausible prepositional relationship and from the desired outcome 1 for a plausible prepositional relationship and still be considered correct. The *error rate* is the percentage of incorrectly classified prepositional relationships in the training set or in the test set. The *total error* is the total sum squared error on the complete training set as defined in (Rumelhart et al. 86, p. 323).

For the training set, the three networks of the three runs showed an error rate between 6.5% and 6.9% using an error tolerance of 0.49, and between 7.3% and 7.7% using an error tolerance of 0.3 (see figure 3). A network which was not trained at all was tested with the training set and showed an error rate of 54.0% for the error tolerance 0.49, and an error rate of 73.4% for the error tolerance 0.3. These tests with the training examples demonstrate that an effective representation for prepositional relationships can be learned.

| Run | 1 | 2 | 3 | No learning |
|---|---|---|---|---|
| Total error at the start of the training | 77.8 | 62.5 | 70.1 | - |
| Total error at the end of the training | 6.6 | 7.1 | 8.1 | - |
| Error rate for the training set for error tolerance 0.49 | 6.9 | 6.5 | 6.9 | 54.0 |
| Error rate for the training set for error tolerance 0.30 | 7.7 | 7.3 | 7.7 | 73.4 |

Figure 3: Test results for the training set for the prepositional relationships of "in"

After the networks had been tested with the 248 training examples, we tested the networks with 30 new test examples which were not part of the training set. For the test set we chose 15 plausible and 15 implausible prepositional relationships from our corpus with the only constraint that the prepositional relationships in the test set were not part of the training set. Examples from the test set are shown in Figure 4.

| Plausible Prepositional Relationships | Implausible Prepositional Relationships |
|---|---|
| Effect in ferromagnetics | Japan in investigation |
| Distortion in amplifier | Power-supply in diode |

Figure 4: Examples of the test set for the prepositional relationships for "in"

The test results with 30 new prepositional relationships showed an error rate between 16.7% and 26.7% for the error tolerance 0.49 and between 20% and 30% for the error tolerance 0.3 (see figure 5). The performance of the trained network on the new test examples can be demonstrated by comparing the described error rates with an untrained network. Training for 1600 epochs reduces the error rate for test examples which were not in the training set from 53.3% to 16.7% for an error tolerance of 0.49, and from 70.0% to 20.0% for an error tolerance of 0.3.

| Run | 1 | 2 | 3 | No learning |
|---|---|---|---|---|
| Error rate for the test set for error tolerance 0.49 | 16.7 | 26.7 | 16.7 | 53.3 |
| Error rate for the test set for error tolerance 0.30 | 20.0 | 30.0 | 20.0 | 70.0 |

Figure 5: Test results for the test set for the prepositional relationships of "in"

To sum up the results for training the backpropagation networks with prepositional relationships for "in", we have shown that for an error tolerance 0.49 trained networks provide the plausibility value of a prepositional relationship correctly in about 93% of the prepositional relationships in the training set and in about 83% of the prepositional relationships in the test set.

## 3.3 Learned Internal Representations for the Prepositional Relationships for "in"

After training had been completed we examined the internal representation in the backpropagation network. Figure 6 illustrates the activation values of the hidden units for 10 training examples. The first 5 rows show the hidden units for training examples with a plausible prepositional relationship, the last 5 rows show the hidden units for training examples with an implausible prepositional relationship.

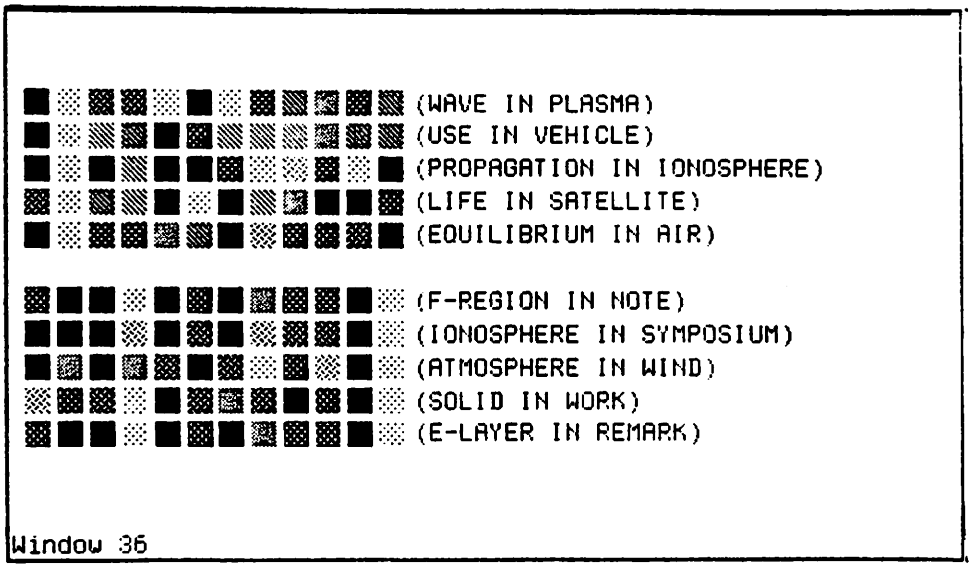


Figure 6: The hidden units for prepositional relationships from the training set

Each row contains the 12 hidden units for one training instance. The hidden units have activation values between 0 (white) and 1 (black). Comparing the internal representations of the plausible prepositional relationships and the implausible prepositional relationships we found that plausible relationships correlate with a low value for hidden unit 2 and a high value for hidden unit 12. Implausible relationships correlate with a high value for hidden unit 2 and a low value for hidden unit 12. We do not claim that these two units are exclusively responsible for the distinction between plausible and implausible prepositional relationships. However, there is a strong tendency for these two units at least to play an important role in the internal representation of this distinction.

After looking at the hidden units with respect to the plausibility of a prepositional relationship we asked if the hidden units represent different word senses for plausible prepositional relationships. We used a simple clustering algorithm to cluster the vectors of 12 hidden units of the plausible prepositional relationships. This clustering algorithm takes a set of prototype vectors as its input and classifies all instances according to the minimal distance to the given prototype vectors. An instance is assigned to the class with the smallest distance to the prototype vector. This distance is computed as the sum of the squared differences between the feature vector of the current instance and the feature vector of the prototype. Although this simple clustering method relies on knowing "good" prototype vectors, this method serves as a first approximation for a classification of the hidden units.

In figure 7 we show examples of the internal representation for 3 clusters. The prototypes for the 3 clusters are the prepositional relationships "effect in rectifier", "radiation in atmosphere", and "effect in beam". These prototypes were chosen because they illustrate different interpretations of "in": "in a physical/electrical object", "in a spatial location/gas", and "in energy", respectively. In comparing the hidden units of these clusters, we found that units 6 and 9 essentially contain the information for differentiating these interpretations. In the first cluster, unit 6 has a low activation value and unit 9 has a high activation value; in the second cluster, unit 6 has a high activation value and unit 9 has a low activation value; and in the last cluster, both units 6 and 9 have low activation values. Although we found a few instances in the training set which use different units to distinguish between these clusters, most prepositional relationships

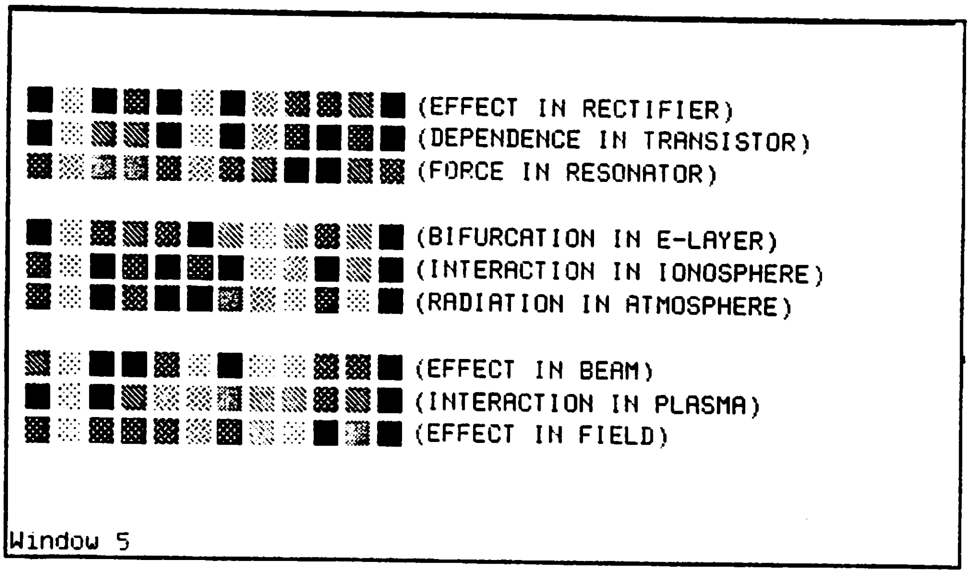


Figure 7: The hidden units for plausible prepositional relationships from 3 clusters

in the three clusters can be differentiated solely based on the units 6 and 9.

To sum up, we have shown that specific hidden units in the distributed internal representation of the learned prepositional relationships are involved in encoding the plausibility of a relationship and in encoding specific interpretations for the prepositional relationship "in". Backpropagation networks which are trained for 1600 epochs can learn effective distributed representations of prepositional relationships. We demonstrated that these network representations currently reach a performance of about 93% (error rate about 7%) on the training set of prepositional relationships and about 83% (error rate about 17%) on the test set of prepositional relationships. Although we describe in detail only the training results for the prepositional relationships for "in", we showed elsewhere (Wermter 1989b) that other prepositional relationships behave very similarly. Experiments with semantic relationships for seven prepositions (by, for, from, in, of, on, with) demonstrate that the results for the prepositional relationships "in" hold for other prepositional relationships as well.

# 4. Integration of Semantic Relationships with Syntactic Constraints in Localist Connectionist Networks

While the last section focused on learning semantic prepositional relationships with backpropagation networks, we now turn to a description of the localist network level. First, we will briefly describe some syntactic constraints in noun phrases. Then, we will show how simple syntactic constraints and learned semantic constraints can be integrated in a localist connectionist network for disambiguating noun phrases.

## 4.1 Syntactic Constraints

The two syntactic constraints we consider are the locality constraint and the no-crossing constraint. The *locality constraint* says that a prepositional phrase is more likely to attach to a close preceding noun than to a distant preceding noun. For instance, in the noun phrase "Techniques for measurements in discharges" the prepositional phrase "in discharges" might attach to "measurements" or to "techniques". The locality constraint suggests that "in discharges" attaches to "measurements" because "measurements" is closer than "techniques".

The *no-crossing constraint* (Tait 1983) for noun phrases means that branches for attachment do not cross. The following (constructed) example shows a violated no-crossing constraint:

Influence of the temperature on the electrons in Fahrenheit

## 4.2 Localist Connectionist Networks for the Integration of Multiple Constraints

Localist networks have been used for a number of tasks to integrate multiple constraints in natural language processing, for instance for sentence understanding (Waltz & Pollack 1985) (Lehnert 87) (Lehnert 88), for word sense disambiguation (Bookman 1987), and for lexical access (Cottrell 1988). We have demonstrated elsewhere that localist networks are useful for integrating semantic and syntactic constraints for noun phrase disambiguation (Wermter 1989a). In this section we describe the most important properties of an efficient localist network that performs noun phrase disambiguation with fewer nodes (see figure 8).

Our localist network consists of three types of nodes: *noun nodes* represent the nouns in a noun phrase, *semantic nodes* represent the plausibility of prepositional relationships between nouns, and *locality nodes* represent the distance between two nouns in a noun phrase. Each node has an activation potential between 0 and 10. A semantic node in the localist network is initialized with the plausibility value of the output unit of

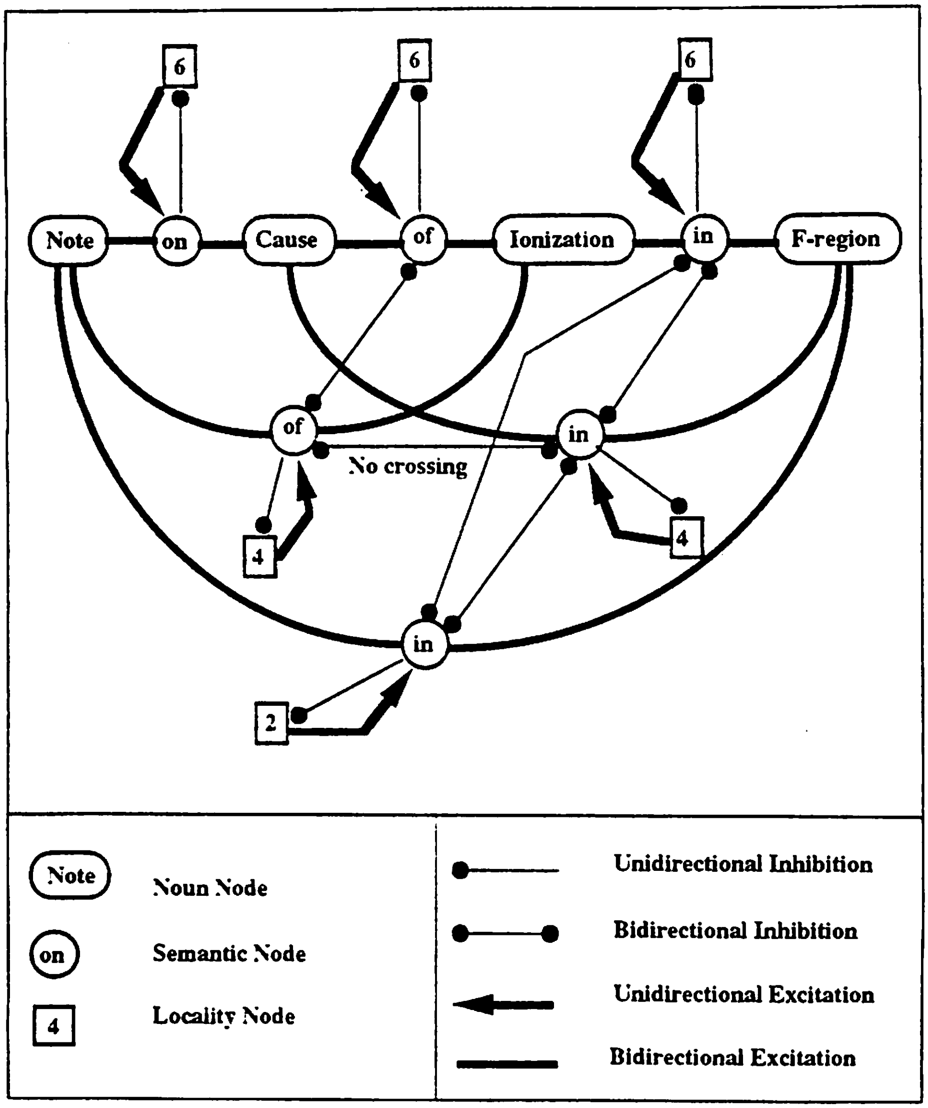


Figure 8: Localist network for the integration of multiple constraints

the appropriate backpropagation network (multiplied by a factor of 10 to get values between 0 and 10). The higher the plausibility value of a prepositional relationship, the higher the initialization value for the semantic node. The initialization of the locality nodes is based on the relative distance between the nouns. The closer two nouns are in a noun phrase, the higher the initialization value for the locality node between these nouns. It is important to point out that the initialization of the locality nodes is fairly independent of specific values (e.g., 6,4, and 2 in figure 8). Other initialization values (e.g., 3,2,1 or 8,4,2) work as well as long as there is a monotonically decreasing relationship for the distances between the nouns, and as long as all the values are not too close to the upper and lower bounds of the nodes. Noun nodes are initialized with 0 activation since they serve only as the framework to which semantic nodes and locality nodes connect. The semantic constraints are encoded as the semantic nodes, the locality constraints as the locality nodes, and the no-crossing constraints as specific inhibitory connections between semantic nodes in crossing attachment links (see figure 8).

All nodes are connected via inhibitory and excitatory connections, as figure 8 shows for a network with three prepositions. Each noun node has excitatory attachment links to each noun node of preceding nouns. The semantic nodes in competing attachment links are inhibitorily connected. The locality nodes provide excitation to semantic nodes depending on the distance of the attachment. The inhibitory connection from a semantic node to a locality node prevents the locality node from sending too much excitation to the semantic node. Networks for noun phrases with a different number of prepositions are built in exactly the same systematic manner. The input nodes are initialized and activation spreads through the network according to a standard relaxation algorithm (Feldman & Ballard 1982). After about 20 to 30 cycles the localist network settles in a global interpretation, and the semantic nodes with the highest activation values determine the preferred structural interpretation. We will see examples of this process and for the interaction of the localist network with the symbolic structures in the next sections.

# 5. Symbolic Level

The last two sections explained the connectionist networks for learning semantic relationships and for integrating semantic and syntactic constraints. In this section we describe the symbolic level of our model for noun phrase understanding.

In the last two sections we have assumed that the noun phrases for the connectionist networks only consist of nouns and prepositions. However, as our examples in section 2 showed, noun phrases often contain other parts of speech as well, including adjectives, adverbs, and determiners. Although these parts of speech might contain significant information they are usually less important for the representation of the *essential* concept of a noun phrase and the structural disambiguation of the noun phrase. Therefore, the purpose of the symbolic level is to extract the essential sequence of nouns and prepositions from the complete noun phrase. Then, this essential reduced noun phrase has the canonical form of nouns and prepositions required for our connectionist levels.

The first mechanism is an analysis of the noun phrase with respect to its syntactic constituents. This restricted syntactic analyis is provided by a subsystem of CIRCUS (Lehnert 1988) which uses a stack-based architecture to recognize simple syntactic constituents. Since we want to extract the essential reduced noun phrase from the complete noun phrase and since we do not need complete parse trees for this extraction, the restricted syntactic analysis in CIRCUS is sufficient for our purpose.

This subsystem uses a syntactic dictionary and syntactic predictions for identifying constituents. The syntactic predictions are encoded as requests and the request packet mechanism of McELI (Schank and Riesbeck 1981) is used to process the predicted next constituents. Before we begin to analyze a noun phrase, an initial syntactic prediction for the head noun will be on top of the stack. This prediction allows us to skip possible intervening constituents like adjectives, adverbs, and determiners and stores the head noun in a global buffer. At this point the current request is removed from the stack. If a preposition follows, then a new request is pushed on the stack for the following prepositional phrase. As soon as the next noun is identified it is stored in another global buffer for this prepositional phrase. This process of adding syntactic predictions, removing the predictions, finding the desired constituents (prepositions and nouns), and storing them in global buffers is continued until the noun phrase is completely processed.

Although it might seem that a simple pattern matching algorithm which identifies nouns and prepositions using a syntactic dictionary might be sufficient, such a simple approach does not account for more complicated noun phrases with associated subclauses or participle constructions. For instance, for the noun phrase "the man in the satellite which blinked in the sun" we only want to extract "man in satellite" and not "man in satellite in sun", which would be constructed in a simple pattern matching approach based solely on parts of speech. Using a syntactic prediction for a new subclause associated with the relative pronoun "which", it is possible to detect and skip this subclause so that only syntactically desired constituents are extracted.

While the restricted syntactic analysis transforms a noun phrase into a reduced noun phrase based on syntactic predictions, a second mechanism can extract the essential part of this reduced noun phrase based on semantic predictions. The semantic predictions are associated with words in the semantic dictionary. Semantic predictions are fulfilled if the current part of speech in the noun phrase is considered essential. The question of what is considered essential depends on the application and the domain. For instance, in an information retrieval context we might have queries like:

Information on papers about turbulences in gas.

In this domain it is not wise to include the nouns "information" and "papers" in a concept representation since they do not contribute any important distinguishing information. For this information retrieval task only the nouns that are important for the domain can fulfill the semantic predictions. In our example these nouns are "turbulence" and "gas", but not "information" and "papers". Therefore "turbulence in gas" is extracted as the essential part of the noun phrase in this application. In general, these semantic predictions are fulfilled if the current constituent is an essential part of the noun phrase or if a preceding constituent was identified as an essential part.

Both semantic and syntactic predictions allow us to extract the essential part of a noun phrase. We will see specific examples of this level in the next section.

## 6. Operation of the System

In this section we describe the operation of our whole system and show some examples of its performance. First, we focus on how a typical noun phrase is processed in detail:

Note on a new cause of increasing ionization in the F-region.

The *symbolic level* extracts the sequence of nouns and prepositions from this noun phrase and provides the following noun phrase:

Note on cause of ionization in F-region.

This symbolic level could also skip relative clauses as in "ionization in the F-region which is close to the Antarctic" or participle constructions as in "ionization in the F-region surrounding the Antarctic". Then all possible prepositional relationships for the reduced noun phrase are computed:

Note on cause

Cause of ionization

Note of ionization

Ionization in F-region

Cause in F-Region

Note in F-region

The feature representation of each noun in the prepositional relationships is looked up in the lexicon. Based on these features, backpropagation networks at the *distributed level* are initialized for each prepositional relationship. The output of the backpropagation networks are the plausibility values for the prepositional relationships. These plausibility values initialize the semantic nodes in the network at the *localist level.* Locality nodes and noun nodes are initialized as well. Then the localist network starts processing, integrates the syntactic and semantic constraints, and stabilizes in a global interpretation of the noun phrase. The activation of the semantic nodes in the localist network determines the preferred structural interpretation of the noun phrase. In our example there are three semantic nodes that have high activation values after the relaxation. These nodes correspond to the following interpretation:

Note on cause of ionization in F-region.

In the following we show more examples of noun phrases and their structural interpretation:

```
(1) Effect of field on turbulence in gas   -->

    CONCEPT: effect
                OF-REL: field
                ON-REL: turbulence
                           IN-REL: gas
```

```
(2) Dependence of amplification in phosphor on intensity   -->

    CONCEPT: dependence
                OF-REL: amplification
                              IN-REL: phosphor
                ON-REL: intensity


(3) Distortion in amplifier on satellite in Van-Allen-belt   -->

    CONCEPT: distortion
                IN-REL: amplifier
                              ON-REL: satellite
                                            IN-REL: Van-Allen-belt


(4) Experiment on diffraction of ray in layer   -->

    CONCEPT: experiment
                ON-REL: diffraction
                              OF-REL: ray
                                           IN-REL: layer
                                                     -----
```

The last example (4) shows that not *all* attachments are necessarily wrong for an interpretation to be considered incorrect. The first two attachments are correct but "in layer" should attach to "diffraction" rather than to "ray". Nevertheless, we consider structural interpretations with at least one wrong attachment to be incorrect. Using this strict and conservative evaluation we tested our system with 80 noun phrases containing up to three prepositions. A correct structural interpretation was assigned for 88% of 50 noun phrases that contained prepositional relationships from the training set and for 77% of 30 noun phrases that contained prepositional relationships which were not in the training set. Prepositional phrases can attach to several nouns if only semantic constraints are considered. The overall strategy is to prefer semantic constraints over syntactic constraints (locality and no-crossing of branches) and to use syntactic constraints to favor one of several possible semantic interpretations.

# 7. Discussion

In this section we first compare our hybrid model with other symbolic models for structural noun phrase disambiguation. Then we focus on the single levels of our model and explain why we chose a hybrid 3-level model.

Recently there has been a lot of interest in attacking the problem of structural ambiguity, especially in prepositional phrase attachment (Wilks et al. 1985) (Schubert 1986) (Dahlgren & McDowell 1986) (Jensen & Binot 1987) (St. John & McClelland 1988). All these approaches focus on attaching a *single* prepositional phrase within a sentence of the form <NP> <VP> <NP> <PP>. Our approach focuses on attaching *multiple* prepositional phrases within noun phrases. Attaching multiple prepositional phrases in noun phrases is a much harder problem since we cannot rely on predictive verbal knowledge alone.

Most previous work on prepositional phrase attachment relies on an intuitive development of symbolic heuristic rules (Wilks et al. 1985) (Schubert 1986) (Dahlgren & McDowell 1986) (Hirst 1987). Since prepositional phrase attachment can not reasonably be attacked without semantic knowledge, these rules have to encode the semantic knowledge and have to be redesigned for new domains. Our model attacks this problem by learning and generalizing over semantic constraints and eliminating knowledge which has to be handcoded.

Another approach for reducing the amount of knowledge engineering can be found in (Jensen & Binot 1987). This approach tackles the problem of acquiring semantic knowledge for attachments by using definitions in an on-line dictionary. Although this symbolic approach was shown to attach correctly a single prepositional phrase in some sentences, this method depends on suitable definitions in the lexicon. While using on-line dictionaries is a very reasonable attempt, it appears that much more work is required in standardizing semantic knowledge in on-line dictionaries before we can use them to support disambiguation in a general manner.

Very recent work on symbolic prepositional phrase attachment (Dahlgren 1988) reports a success rate above 93% for the attachment of *single* prepositional phrases. These results were obtained by hand-testing intuitively developed rules on several small corpora. Our approach reaches 88% on the training set and 77% on the test set of new noun phrases. Although our results might be even better with further training we believe that our current results already demonstrate the effectiveness of our approach for two reasons. First, multiple prepositional phrase attachment is a much harder problem than Dahlgren's single prepositional phrase attachment. In our experiments we considered noun phrases with up to three prepositional phrases. Second, our model did not rely on intuitively developed rules but learned part of its knowledge. In general, we believe that our hybrid model has a lot of potential compared with traditional purely symbolic methods since our hybrid model attacks a much harder problem, acquires part of its knowledge by learning, and already comes close to the best performance of purely symbolic approaches that attack a significantly simpler problem.

We now turn to the discussion of the three levels in our hybrid model and give reasons for the design of each individual level. The *symbolic level* performs a restricted syntactic analysis and extracts the essential concept (the sequence of nouns and prepositions) of a noun phrase for the attachment decision. A symbolic approach is more suitable for this level since the extraction of the essential concept based on syntactic and semantic predictions is a sequential control problem – it has to decide which constituents to process. In a symbolic mechanism, syntactic and semantic predictions for this extraction can be formulated easily. In contrast, in a connectionist framework localist networks would have to be designed or distributed networks trained to perform the extraction. Although there has been some success using recurrent networks for processing restricted sequential structures (Pollack 1988) (St. John & McClelland 1988) (Elman 1988), these recurrent networks do not seem to be powerful enough for this high level sequential control problem. Since symbolic techniques are particularly suitable for dealing with sequentiality and control, they are more useful for dealing with the variety of constituents as they occur in noun phrases in real world examples.

The *localist level* performs the integration of syntactic and semantic constraints. The links in the localist network implement the possible attachments and their mutual competition. The localist model considered here is more efficient in the number of nodes than the localist model for attachment in (Wermter 1989a). That model used one structure node for each possible structural interpretation of a noun phrase. While that architecture was useful for short noun phrases, the number of structure nodes increased exponentially with the length of the noun phrase. In our present model the total number of nodes in the network increases only quadratically with the length of the noun phrases.

A similar localist network for prepositional phrase attachment in noun phrases can be found in (Touretzky 1989). While Touretzky uses a similar attachment architecture, he implements locality constraints only in a restricted way by reducing the specific threshold for the unit representing the "nearest neighbor" noun. In our model we implemented locality constraints explicitly in a more general way with locality nodes for the relationship to every preceding noun. For the locality nodes, we found empirically that initialization values should decrease monotonically with the length of the attachment and they should be well under the upper threshold for the nodes[3].

The initialization of the semantic nodes is based on the *distributed level.* The dis-

[3]For example, for a noun phrase with three prepositions initialization values of 3,2,1 for the different attachments implement a small syntactic locality effect. The values 6,4,2 implement a moderate syntactic influence. If the values are too high, e.g., 10,9,8, the network gets overloaded with excitation from the locality nodes.

tributed backpropagation networks learn and generalize the semantic prepositional relationships and provide a semantic memory model for the initialization of the localist network. Other work on learning relationships between constituents (Hinton 1986) (Cosic & Munro 1988) can not be directly used to provide this memory model. Hinton attacks a completion task that finds a specific relative given a person and a family relationship. Cosic and Munro tackle a completion task which determines the meaning of a preposition based on the lexical item of the preposition and two nouns. Although this work deals with learning relationships, these architectures can not be directly used to support the initialization of single nodes in localist networks.

Furthermore, both architectures (Hinton 1986) (Cosic & Munro 1988) have all constituents and all relationships encoded in one backpropagation network. While this might be sufficient for small applications, one huge network can not be expected to be efficient in terms of training time and generalization behavior for scaling up to bigger applications. Therefore, we have one backpropagation network for each preposition. Apart from less training time and better generalization, this modular architecture also allows us the modification and addition of individual prepositions without retraining the whole network.

Another interesting design issue is the number of units in the backpropagation networks. The number of input units was determined by our choice of 16 features for representing each noun. There is one output unit for the plausibility value. More interesting is our choice of 12 for the number of hidden units. Increasing the number of 12 hidden units led to better performance on the training set, but worse performance on the test set.

Decreasing the number of hidden units decreased the performance on the training set and test set. Apparently, there is a tradeoff between memorization and generalization and we found our best results with a hidden layer that had slightly less than half the number of input units.

# 8. Conclusion

We have described a hybrid symbolic/connectionist system for noun phrase disambiguation. The symbolic level supplies input for the connectionist networks by extracting the sequence of nouns and prepositions from a noun phrase. The localist connectionist network integrates semantic and syntactic constraints for noun phrase disambiguation and computes a preferred structural interpretation. Distributed connectionist networks learn semantic relationships between nouns, allow for generalizations of the learned relationships, and provide a semantic memory for initializing nodes in the localist connectionist networks. This hybrid three-level model of distributed connectionist networks, localist connectionist networks, and symbolic concepts allows for the combination of learning and generalization, the integration of competing constraints, and the symbolic extraction of concepts and makes this hybrid model potentially stronger than models relying on techniques from only one of the three processing paradigms.

# Acknowledgements

We would like to thank Claire Cardie, David Lewis, Ellen Riloff, and the anonymous reviewers for comments on earlier drafts of this paper. This research was supported by the Advanced Research Projects Agency of the Department of Defense, monitored by the Office of Naval Research under contract #N00014-87-K-0238, the Office of Naval Research, under a University Research Initiative Grant, Contract #N00014-86-K-0764 and NSF Presidential Young Investigators Award NSFIST-8351863.

## References

Allen R.B. (1987) Several studies on natural language and back-propagation. Proceedings of the First Annual International Conference on Neural Networks.

Bookman L.A. (1987) A microfeature based scheme for modelling semantics. Proceedings of the International Joint Conference on Artificial Intelligence.

Charniak E. (1983) A parser with something for everyone. In: King M. (Ed.) Parsing Natural Language. Academic Press, London.

Cosic C., Munro P. (1988) Learning to represent and understand locative prepositional phrases. Proceedings of the Tenth Annual Conference of the Cognitive Science Society.

Cottrell G.W., Small S.L. (1983) A connectionist scheme for modelling word sense disambiguation. Cognition and Brain Theory 6.

Cottrell G.W. (1988) A Model of Lexical Access of ambiguous words. In: Small S.I., Cottrell G.W., Tanenhaus M.K. (Eds.) Lexical Ambiguity Resolution. Morgan Kaufmann, San Mateo.

Crain S., Steedman M. (1985) On not being led up the garden path: The use of context by the psychological syntax processor. In: Dowty D.R., Karttunen L., Zwicky A.M. (Eds.)

Natural Language Parsing: Psychological, computational, and theoretical perspectives. Cambridge University Press, Cambridge.

Dahlgren K., McDowell J. (1986) Using Commonsense Knowledge to disambiguate Prepositional Phrase Modifiers. Proceedings of the National Conference on Artificial Intelligence.

Dahlgren K. (1988) Naive Semantics for Natural Language Understanding. Kluwer Academic Publishers.

Derthick M. (1988) Mundane Reasoning by parallel constraint satisfaction. Technical Report, CMU-CS-88-182, Carnegie-Mellon University.

Dolan C.P., Dyer M.G. (1988) Parallel Retrieval and Application of Conceptual Knowledge. Technical Report UCLA-AI-88-3, University of California, Los Angeles.

Dyer M.G. (1983) In-depth understanding. MIT Press, Cambridge.

Dyer M.G. (1988) Symbolic NeuroEngineering for Natural Language Processing: A multilevel research approach. In: Barnden J., Pollack J. (Eds.) Advances in Connectionist and Neural Computation Theory, Ablex Publishers, in press (also available as Technical Report UCLA-AI-88-14, University of California, Los Angeles).

EJC (1967) Thesaurus of Engineering and Scientific Terms. Engineers Joint Council, New York.

Elman J.L. (1988) Finding Structure in Time. Technical Report CRL-8801, University of California, San Diego.

Fanty, M. (1985) Context-free parsing in connectionist networks. Technical Report 174, University of Rochester, Rochester.

Feldman J.A., Ballard D.H. (1982) Connectionist Models and their properties. Cognitive Science 6.

Ford M., Bresnan J.W., Kaplan R.M. (1982) A competence based theory of syntactic closure. In: Bresnan J.W. (Ed.) The mental representation of grammatical relations.

Frazier L., Fodor J.D. (1978) The sausage machine: A new two-stage parsing model. Cognition 6.

Gasser M.E. (1988) A connectionist model of sentence generation in a first and second language. Technical Report UCLA-AI-88-13, University of California, Los Angeles.

Grosz B.J., Appelt D.E., Martin P.A., Pereira F.C.N. (1987) TEAM: An Experiment in the Design of Transportable Natural-Language Interfaces. Artificial Intelligence 32.

Hanson S.J., Kegl J. (1987) PARSNIP: A connectionist network that learns natural language grammar from exposure to natural language sentences. Proceedings of the Ninth Annual Conference of the Cognitive Science Society.

Hendler, J.A. (1989) Marker Passing over Microfeatures: Towards a hybrid symbolic/connectionist model. Cognitive Science 13.

Hinton G.E. (1986) Learning Distributed Representations of Concepts. Proceedings of the Eighth Annual Conference of the Cognitive Science Society.

Hinton G.E. (1988) Representing part-whole hierarchies in connectionist networks. Proceedings of the Tenth Annual Conference of the Cognitive Science Society.

Hirschman L. (1986) Discovering Sublanguage Structures. In: Grishman R., Kittredge, R. (Eds.) Analyzing Language in Restricted Domains: Sublanguage Description and Processing.

Hirst G. (1987) Semantic interpretation and the resolution of ambiguity. Cambridge University Press, Cambridge.

Howells, T. (1988) VITAL - A connectionist parser. Proceedings of the Tenth Annual Conference of the Cognitive Science Society.

Jensen K., Binot J.L. (1987) Disambiguating Prepositional Phrase Attachments by Using Online Dictionary Definitions. Computational Linguistics 13.

Jordan, M.I. (1986) Attractor dynamics and parallelism in a connectionist sequential machine. Proceedings of the 8th Annual Conference of the Cognitive Science Society.

Kimball J. (1973) Seven principles of surface structure parsing in natural language. Cognition 2.

Kwasny S.C. (1988) A parallel distributed approach to parsing natural language deterministically. Technical Report WUCS-88-21, Washington University, St. Louis.

Lehnert W.G. (1987) Learning to integrate syntax and semantics. Proceedings of the Fourth International Workshop on Machine Learning.

Lehnert W.G. (1988) Symbolic/Subsymbolic Sentence Analysis: Exploiting the Best of Two

Worlds. In: Barnden J., Pollack J. (Eds.) Advances in Connectionist and Neural Computation Theory, Ablex Publishers, in press (also available as COINS Technical Report 88-99, Department of Computer and Information Science, University of Massachusetts, Amherst, MA).

NASA (1985) NASA Thesaurus National Aeronautics and Space Administration.

McClelland J.L., Kawamoto A.H. (1986) Mechanisms of Sentence Processing: Assigning Roles to Constituents. In: McClelland J.L., Rumelhart D.E. (Eds.) Parallel Distributed Processing Vol. 2. MIT Press, Cambridge.

Miikulainen R., Dyer M.G. (1989) A modular neural network architecture for sequential paraphrasing of script-based stories. Technical Report UCLA-AI-89-02, University of California, Los Angeles.

Pollack J. (1988) Recursive Auto-Associative Memory: Devising Compositional Distributed Representations. Proceedings of the Tenth Annual Conference of the Cognitive Science Society.

Pollack J. (1989) Implications of Recursive Distributed Representations. Submitted to the Conference on Neural Information Processing Systems.

Riesbeck C., Martin C. (1986) Direct Memory Access Parsing. In: Riesbeck C., Kolodner J. (Eds.) Experience, Memory and Reasoning.

Rumelhart D.E., Hinton G.E., Williams R.J. (1986) Learning Internal Representations by Error Propagation. In: McClelland J.L., Rumelhart D.E. (Eds.) Parallel Distributed Processing Vol 1. MIT Press, Cambridge.

Rumelhart D.E., McClelland J.L. (1986) On learning the past tenses of English verbs. In: McClelland J.L., Rumelhart D.E. (Eds.) Parallel Distributed Processing Vol 1. MIT Press, Cambridge.

Schank R.D., Riesbeck C. (1981) Inside Computer Understanding. Lawrence Erlbaum.

Schubert L.K. (1986) Are there preference tradeoffs in Attachment decisions? Proceedings of the National Conference on Artificial Intelligence.

Sejnowski T.J., Rosenberg C.R. (1986) NETtalk: A parallel network that learns to read aloud. Technical Report JHU/EECS-86/01, Johns Hopkins University, Baltimore.

Shastri L. (1988) A connectionist approach to knowledge representation and limited inference. Cognitive Science 12.

Shastri L., Ajjanagadde V. (1989) A connectionist system for rule based reasoning with multi-place predicates and variables. Technical Report, MS-CIS-89-06, University of Pennsylvania.

Sparck Jones K., VanRijsbergen C.J. (1976) Information Retrieval Test Collections. Journal of Documentation 32.

St. John M.F., McClelland J.L. (1988) Applying contextual constraints in sentence comprehension. Proceedings of the Tenth Annual Conference of the Cognitive Science Society.

Tait J.I. (1983) Semantic parsing and syntactic constraints (Mark IV). In: Sparck Jones K., Wilks Y. (Eds.) Automated natural language processing. Ellis Horwood Limited.

Touretzky D.S., Hinton G.E. (1985) Symbols among the neurons: Details of a connectionist inference architecture. Proceedings of the International Joint Conference on Artificial Intelligence.

Touretzky, D.S., Hinton G.E. (1988) A distributed connectionist production system. Cognitive Science 12.

Touretzky, D.S. (1988) Beyond Associative Memory: Connectionists must search for other cognitive primitives. Proceedings of the AAAI Spring Symposium Series: Parallel Models of Intelligence: How can slow components think so fast?

Touretzky D.S. (1989) Connectionism and Compositional Semantics. Technical Report, CMU-CS-89-147, Carnegie Mellon University.

Waltz D.L., Pollack J.B. (1985) Massively Parallel Parsing: A strongly interactive model of natural language interpretation. Cognitive Science 9.

Wermter S. (1989a) Integration of Semantic and Syntactic Constraints for Structural Noun Phrase Disambiguation. Proceedings of the International Joint Conference on Artificial Intelligence.

Wermter S. (1989b) Learning Semantic Relationships in Compound Nouns with Connectionist Networks. Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.

Wilks Y., Huang X., Fass D. (1985) Syntax, Preference and Right Attachment. Proceedings of the International Joint Conference on Artificial Intelligence.