

GRAMMAR VIEWED AS A FUNCTIONING PART OF A COGNITIVE SYSTEM

Helen M. Gigley
Department of Computer Science
University of New Hampshire
Durham, NH 03824

ABSTRACT

How can grammar be viewed as a functional part of a cognitive system? Given a neural basis for the processing control paradigm of language performance, what roles does "grammar" play? Is there evidence to suggest that grammatical processing can be independent from other aspects of language processing?

This paper will focus on these issues and suggest answers within the context of one computational solution. The example model of sentence comprehension, HOPE, is intended to demonstrate both representational considerations for a grammar within such a system as well as to illustrate that by interpreting a grammar as a feedback control mechanism of a "neural-like" process, additional insights into language processing can be obtained.

1. Introduction

The role of grammar in defining cognitive models that are neurally plausible and psychologically valid will be the focus of this paper. While linguistic theory greatly influences the actual representation that is included in any such model, there are vast differences in how any grammar selected is "processed" within a "natural computation" paradigm. The processing does not grow trees explicitly; it does not transform trees explicitly; nor does it move constituents.

In this type of model, a grammar is an explicit encoded representation that coordinates the integrated parallel process. It provides the interfaces between parallel processes that can be interpreted within semantic and syntactic levels separately. It furthermore acts as a "conductor" of a time-synchronized process. Aspects of how a grammar might be processed within a cognitive view of sentence comprehension will be demonstrated within an implemented model of such processing, HOPE (Gigley, 1981; 1982a; 1982b; 1983; 1984; 1985). This view of grammatical "process" suggests that neural processing should be included as a basis for defining what is universal in language.

2. Background

There are currently several approaches to developing cognitive models of linguistic function (Cottrell, 1984; Cottrell and Small, 1983; Gigley, 1981; 1982a; 1982b; 1983; 1984; 1985; Small, Cottrell and Shastri, 1982; Waltz and Pollack, in press). These models include assumptions about memory processing within a spreading activation framework (Collins and Loftus, 1975; Hinton, 1981; Quillian, 1968/1980), and a parallel, interactive control paradigm for the processing. They differ in the explicit implementations of these theories and the degree to which they claim to be psychologically valid.

Computational Neurolinguistics (CN), first suggested as a problem domain by Arbib and Caplan (1979), is an Artificial Intelligence (AI) approach to modelling neural processes which subserve natural language performance. As CN has developed, such models are highly constrained by behavioral evidence, both normal and pathological. CN provides a framework for defining cognitive models of natural language performance of behavior that includes claims of validity at two levels, the natural computation or neural-like processing level, and at the system result or behavioral level.

Using one implementation of a CN model, HOPE (Gigley, 1981; 1982a; 1982b; 1983) a model of single sentence comprehension, the remainder of the paper will illustrate how the role of grammar can be integrated into the design of such a model. It will emphasize the importance of the parallel control assumptions in constraining the representation in which the grammar is encoded. It will demonstrate how the grammar contributes to control the coordination of the parallel, asynchronous processes included in the model.

The HOPE model is chosen explicitly because the underlying assumptions in its design are intended to be psychologically valid on two levels, while the other referenced models do not make such claims. The complete model is discussed in Gigley (1982a; 1982b; 1983) and will be summarized here to illustrate the role of the grammar in its function. The suggested implications and goals for including neurophysiological evidence in designing such models have been discussed else-

where in Lavorel and Gigley (1983) and will be included only as they relate to the role and function of the grammar.

2.1 Summary of Included Knowledge and its Representation

Types of representations included in the HOPE model, phonetic, categorially accessed meanings, grammar, and pragmatic or local context, receive support as separately definable knowledge within studies of aphasia. There is a vast literature concerning what aspects of language are independently affected in aphasia that has been used as a basis for deciding these representations. (See Gigley, 1982b for complete documentation.)

Information that is defined within the HOPE model is presented at a phonological level as phonetic representations of words (a stub for a similar interactive process underlying word recognition). Information at the word meaning level is represented as multiple representations, each of which has a designated syntactic category type and orthographic spelling associate to represent the phonetic word's meaning (also a stub). The grammatical representation has two components. One is strictly a local representation of the grammatical structural co-occurrences in normal language. The other is a functional representation, related to interpretation, that is unique for each syntactic category type. Please note that type is not used in the strictest sense of its use within a typed semantic system. This will be described in detail later. Finally, the pragmatic interpretation is assumed to reflect the sentential context of the utterance.

Each piece of information is a thresholding device with memory. Associational interconnections are made by using an hierarchical graph which includes a hypergraph facility that permits simultaneous multiple interpretations for any active information in the process. Using this concept, an active node can be ambiguous, representing information that is shared among many interpretations. Sentence comprehension is viewed as the resolution of the ambiguities that are activated over the time course of the process.

Within our implementation, graphs can represent an aspect of the problem representation by name. Any name can be attached to a node, or an edge, or a space (hypergraph) of the graph. There are some naming constraints required due to the graph processing system implementation, but they do not affect the conceptual representation on which the encoding of the cognitive linguistic knowledge relies.

Any name can have multiple meanings associated with it. These meanings can be interpreted differently by viewing each space in which the name is referenced as a different viewpoint for the same information. This means that whenever the name is the same for any information, it is indeed the same information, although it may mean several things simultaneously. An example related to the grammatical representation is that the

syntactic category aspect of each meaning of a phonetic word is also a part of the grammatical representation where it makes associations with other syntactic categories. The associations visible in the grammatical representation and interpreted as grammatical "meanings" are not viewable within the phonetic word meaning perspective.

However, any information associated with a name, for instance, an activity value, is viewable from any spaces in which the name exists. This means that any interpreted meaning associated with a name can only be evaluated within the context, or contexts, in which the name occurs. Meaning for any name is contextually evaluable. The explicit meaning within any space depends on the rest of the state of the space, which furthermore depends on what previous processing has occurred to affect the state of that space.

2.2 Summary of the Processing Paradigm

The development of CN models emphasizes process. A primary assumption of this approach is that neural-like computations must be included in models which attempt to simulate any cognitive behavior (Cf Lavorel and Gigley, 1983), specifically natural language processing in this case. Furthermore, CN includes the assumption that time is a critical factor in neural processing mechanisms and that it can be a significant factor in language behavior in its degraded or "lesioned" state.

Simulation of a process paradigm for natural language comprehension in HOPE is achieved by incorporating a neurally plausible control that is internal to the processing mechanism. There is no external process that decides which path or process to execute next based on the current state of the solution space. The process is time-locked; at each process time interval. There are six types of serial-order computations that can occur. They apply to all representation viewpoints or spaces simultaneously, and uniformly. Threshold firing can affect multiple spaces, and has a local effect within the space of firing.

Each of these serial-order computations is intended to represent an aspect of "natural computation" as defined in Lavorel and Gigley, 1983. A natural computation, as opposed to a mechanistic one, is a "computation" that is achieved by neural processing components, such as threshold devices and energy transducers, rather than by components such as are found in digital devices. The most important aspect of the control is that all of the serial order computations can occur simultaneously and can affect any information that has been defined in the instantiated model.

Processing control is achieved using activity values on information. As there is no preset context in the current implementation, all information initially has a resting activity value. This activity value can be modified over time depending on the sentential input. Furthermore, there is an automatic activity decay scheme intended to represent memory processing which is

based on the state of the information, whether it has reached threshold and fired or not.

Activity is propagated in a fixed-time scheme to all "connected" aspects of the meaning of the words by spreading activation (Collins and Loftus, 1975; 1983; Hinton, 1981; Quillian, 1968/1980). Simultaneously, information interacts asynchronously due to threshold firing. A state of threshold firing is realized as a result of summed inputs over time that are the result of the fixed-time spreading activation, other threshold firing or memory decay effects in combination. The time course of new information introduction, which initiates activity spread and automatic memory decay is parameterized due to the underlying reason for designing such models (Gigley, 1982b; 1983; 1985).

The exact serial-order processes that occur at any time-slice of the process depend on the "current state" of the global information; they are context dependent. The serial-order processes include:

- (1) NEW-WORD-RECOGNITION: Introduction of the next phonetically recognized word in the sentence.
- (2) DECAY: Automatic memory decay exponentially reduces the activity of all active information that does not receive additional input. It is an important part of the neural processes that occur during memory processing.
- (3) REFRACTORY-STATE-ACTIVATION: An automatic change of state that occurs after active information has reached threshold and fired. In this state, the information can not affect or be affected by other information in the system.
- (4) POST-REFRACTORY-STATE-ACTIVATION: An automatic change of state which all fired information enters after it has existed in the REFRACTORY-STATE. The decay rate is different than before firing, although still exponential.
- (5) MEANING-PROPAGATION: Fixed-time spreading activation to the distributed parts of recognized words' meanings.
- (6) FIRING-INFORMATION-PROPAGATION: Asynchronous activity propagation that occurs when information reaches threshold and fires. It can be INHIBITORY and EXCITATORY in its effect. INTERPRETATION results in activation of a pragmatic representation of a disambiguated word meaning.

Processes (2) through (6) are applicable to all active information in the global representation, while process (1) provides the interface with the external input of the sentence to be understood. The state of the grammar representation affects inhibitory and excitatory firing propagation, as well as coordinates "meaning" interpretation with on-going "input" processing.

It is in the interaction of the results of these asynchronous processes that the process of comprehension is simulated.

3. The Role of a Grammar in Cognitive Processing Models

Within our behavioral approach to studying natural language processing, several considerations must be met. Justification must be made for separate representations of information and, whenever possible, neural processing support must be found.

3.1 Evidence for a Separate Representation of Grammar

Neurolinguistic and psycholinguistic evidence supports a separately interpretable representation for a grammar. The neurolinguistic literature demonstrates that the grammar can be affected in isolation from other aspects of language function. (Cf Studies of agrammatic and Broca's aphasia as described in Goodenough, Zurif, and Weintraub, 1977; Goodglass, 1976; Goodglass and Berko, 1960; Goodglass, Gleason, Bernholtz, and Hyde, 1970; Zurif and Blumstein, 1978).

In the HOPE model, this separation is achieved by including all relevant grammatical information within a space or hypergraph called the grammar. The associated interpretation functions for each grammatical type provide the interface with the pragmatic representation. Before describing the nature of the local representation of the currently included grammar, a brief discussion of the structure of the grammar and the role of the grammar in the global nature of the control must be given.

3.2 The Local Representation of the Grammar

The grammar space contains the locally defined grammar for the process. The current model defined within the HOPE system includes a form of a Categorical Grammar (Ajdukiewicz, 1935; Lewis, 1972). Although the original use of the grammar is not needed, the relationship that ensues between a well defined syntactic form and a "final state" meaning representation is borrowed. Validity of the "final state" meaning is not the issue. Final state here means, at the end of the process. As previously mentioned, typed semantics is also not rigidly enforced in the current model.

HOPE allows one to define a lexicon within user selected syntactic types, and allows one to define a suitable grammar of the selected types in the prescribed form as well. The grammar may be defined to suit the aspects of language performance being modelled.

There are two parts to the grammatical aspect of the HOPE model. One is a form of the structural co-occurrences that constitute context free phrase structure representations of grammar. However, these specifications only make one "constituent" predictions for subsequent input types where each constituent may have additional sub-structure.

Predictions at this time do not spread to substructures because of the "time" factor between computational updates that is used. A spread to substructures will require a refinement in time-sequence specifications.

The second aspect of the representation is an interpretation function, for each specified syntactic type in the grammar definition. Each interpretation function is activated when a word meaning fires for whatever reason. The interpretation function represents a firing activation level for the "concept" of the meaning and includes its syntactic form. For this reason, each syntactic form has a unique functional description that uses the instantiated meaning that is firing (presently, the spelling notation) to activate structures and relations in the pragmatic space that represent the "meaning understood."

Each function activates different types of structures and relations, some of which depend on prior activation of other types to complete the process correctly. These functions can trigger semantic feature checks and morphological matches where appropriate.

Syntactic types in the HOPE system are of two forms, lexical and derived. A lexical category type is one which can be a category type of a lexical item. A derived category type is one which is "composed". Derived category types represent the occurrence of proper "meaning" interpretation in the pragmatic space.

The current represented grammar in HOPE contains the following lexical categories: DET for determiner, ENDCONT for end of sentence intonation, NOUN for common noun, PAUSE for end of clause intonation, TERM for proper nouns, VIP for intransitive verb, VTP for transitive verb. As is seen, the lexical "categories" relate "grammatical" structure to aspects of the input signal, hence in this sense ENDCONT and PAUSE are categories.

The derived categories in the current instantiated model include: SENTENCE, representing a composition of agent determination of a TERM for an appropriate verb phrase, TERM, representing a composed designated DET NOUN referent, and VIP, representing the state of proper composition of a TERM object with a VTP, transitive verb sense. TERM and VIP are examples of category types in this model that are both lexical and derived.

"Rules" in the independently represented grammar are intended to represent what is considered in HOPE as the "syntactic meaning" of the respective category. They are expressed as local interactions, not global ones. Global effects of grammar, the concern of many rule based systems, can only be studied as the result of the time sequenced processing of an "input". Table 1 contains examples of "rules" in our current model. Other categories may be defined; other lexical items defined; other interpretations defined within the HOPE paradigm.

Table 1: Category specification

DET: = TERM / NOUN
 VIP: = SENTENCE / ENDCOUNT
 VTP: = VIP / TERM

In Table 1, the "numerator" of the specification is the derived type which results from composition of the "denominator" type interpretation with the interpretation of the category whose meaning is being defined. For example, DETERMINER, the defined category, combines with a NOUN category type to produce an interpretation which is a TERM type. When a category occurs in more than one place, any interpretation and resultant activity propagation of the correct type may affect any "rule" in which it appears. Effects are in parallel and simultaneous. Interpretation can be blocked for composition by unsuccessful matches on designated attribute features or morphological inconsistencies.

Successful completion of function execution results in a pragmatic representation that will either fire immediately if it is non-compositional or in one time delay if the "meaning" is composed. Firing is of the syntactic type that represents the correctly "understood" entity. This "top-down" firing produces feedback activity whose effect is "directed" by the state of the grammar, space, i.e. what information is active and its degree of activity.

The nature of the research in its present state has not addressed the generality of the linguistic structures it can process. This is left to future work. The concentration at this time is on initial validation of model produced simulation results before any additional effort on expansion is undertaken. With so many assumptions included in the design of such models, initial assessment of the model's performance was felt to be more critical than its immediate expansion along any of the possible dimensions previously noted as stubs. The initial investigation is also intended to suggest how to expand these stubs.

3.3 The Grammar as a Feedback Control System

The role of the grammar as it is encoded in HOPE is to function in a systems theoretic manner. It provides the representation of the feedforward, or prediction, and feedback, or confirmation interconnections among syntactic entities which have produced appropriate entities as pragmatic interpretations. It coordinates the serial ordered expectations, with what actually occurs in the input signal, with any suitable meaning interpretations that can affect the state of the process in a top-down sense. It represents the interface between the serial-order input and the parallel functioning system.

Grammatical categories are activated via spreading activation that is the result of word meaning activation as words are recognized. Firing of an instance of a grammatical type activates that type's interpretation function which

results in the appropriate pragmatic interpretation for it, including the specific meaning that was fired.

Interpretation functions are defined for syntactic types not specific items within each type. Each type interpretation has one form with specific lexical "parameters". All nouns are interpreted the same; all intransitive verbs the same. What differs in interpretation is the attributes that occur for the lexical item being interpreted. These also affect the interpretation.

The meaning representation for all instances of a certain category have the same meta-structure. General nouns (NOUN) are presently depicted as nodes in the pragmatic space. The node name is the "noun meaning." For transitive verbs, nodes named as the verb stem are produced with a directed edge designating the appropriate TERM category as agent. The effect of firing of a grammatical category can trigger feature propagations or morphological checks depending on which category fires and the current pragmatic state of the on-going interpretation.

Successful interpretation results in threshold firing of the "meaning." This "meaning" has a syntactic component which can affect grammatical representations that have an activity value. This process is time constrained depending on whether the syntactic type of the interpretation is lexical or derived.

3.4 Spreading Activation of the Grammar

Input to HOPE is time-sequenced, as phonetically recognized words, (a stub for future development). Each phonetic "word" activates all of its associated meanings. (HOPE uses homophones to access meanings.) Using spreading activation, the syntactic category aspect of each meaning in turn activates the category's meaning in the grammar space representation.

Part of the grammatical meaning of any syntactic category is the meaning category that is expected to follow it in the input. The other part of the grammatical meaning for any category type, is the type it can derive by its correct interpretation within the context of a sentence. Because each of these predictions and interpretations are encoded locally, one can observe interactions among the global "rules" of the grammar during the processing. This is one of the motivating factors for designing the neurally motivated model, as it provides insights into how processing deviations can produce degraded language performance.

3.5 Grammar State and Its Effect on Processing

Lexical category types have different effects than derived ones with respect to timing and pragmatic interpretation. However, both lexical and derived category types have the same effect on the subsequent input. This section will describe the currently represented grammar and provide

example processing effects that arise due to its interactive activation.

Through spreading activation, the state of the syntactic types represented in the grammar affects subsequent category biases in the input (feedforward) and on-going interpretation or disambiguation of previously "heard" words (feedback). The order of processing of the input appears to be both right to left and left to right. Furthermore, each syntactic type, on firing, triggers the interpretation function that is particular to each syntactic type.




Rules, as previously discussed, are activated during processing via spreading activation. Each recognized word activates all "meanings" in parallel. Each "meaning" contains a syntactic type. Spreading activation along "syntactic type associates" (defined in the grammar) predictively activates the "expected" subsequent categories in the input.

In the HOPE model, spreading activation currently propagates this activity which is not at the "threshold" level. Propagated activity due to firing is always a parameter controlled percentage of the above threshold activity and in the presently "tuned" simulations always propagates a value that is under threshold by a substantial amount.

All activations occur in parallel and affect subsequent "meaning" activities of later words in the sentence. In addition, when composition succeeds (or pragmatic interpretation is finalized) the state of the grammar is affected to produce or changes in category aspects of all active meanings in the process.

The remainder of this section will present instances of the feedforward and feedback effects of the grammar during simulation runs to illustrate the role of grammar in the process. The last example will illustrate how a change in state of the grammar representation can affect the process. All examples will use snapshots of the sentence: "The boy saw the building." This is input phonetically as: (TH-UH B-OY S-AO TH-UH B-IH-L-D-IH-NG).

3.5.1 An Example of Feedforward, Feedback, and Composition

This example will illustrate the feedforward activation of NOUN for the DETERMINER grammatical meaning during interpretation of the initial TERM or noun phrase of the sentence. All figures are labelled to correspond with the text. Each interval is labelled at the top, t1, t2, etc. The size of each node reflects the activity level, larger means more active. Threshold firing is represented as . Other changes of state that affect memory are denoted by  and  and are shown for completeness. They indicate serial-order changes of state described earlier, but are not critical to the following discussion.




Figure 1

On "hearing" /TH-UH/ (a) at t1, the represented meaning "DET-the" is activated as the only meaning (b). At the next time interval, t2, the meaning of DET is activated - which spreads activity to what DET predicts, a NOUN (c). All NOUN meanings are activated by spread in the next time interval, t3, in combination with new activity. This produces a threshold which "fires" the "meaning" selected (d). At completion of interpretation (e), in t4, feedback occurs to all instances of meaning with category types in the grammar associated as predictors of the interpreted category. DET is the only active category that predicts NOUN so all active meanings of type DET will receive the feedback activity. In Figure 1, DET-the is ready to fire (f). The increase or decrease in activity of all related types, competitive ones for the meaning (inhibitory) (g) as well as syntactic ones for composition (excitatory) (f) is propagated at the next interval after firing, shown in t3 and t4. In t5, /S-AO/ enters the process (h) with its associated meanings.

The effect of DET-the firing is also seen in t5 where the compositional TERM is activated (i). NOTE: DETERminers are not physically represented as entities in the pragmatic space. Their meaning is only functional and has a "semantic" compositional effect. Here 'tne' requires a "one and only one" NOUN that is unattached as a TERM to successfully denote the meaning of the boy as a proper TERM (i). As this is a compositional "meaning", the firing will affect t6. Because

there is no active TERM prediction in the grammar space, and no competitive meanings, the top-down effect in t6 will be null and is not shown. The next example will illustrate a top-down effect following TERM composition.

3.5.2 An Example of Feedforward, Feedback, Composition, and Subsequent Feedback

This example, shown in Figure 2, will be very similar to the previous one. Only active information discussed is shown as otherwise the figures become cluttered. The grammar is in a different state in Figure 2 when successful TERM interpretation occurs at t11 (a). This is due to the activation at t9 of all meanings of B-UI-L-O-IH-NG (b).

The VTP meanings of /S-AO/ and then /B-UI-L-O-IH-NG/ make a TERM prediction shown as it remains in t10 (c). After composition of "the building" (a) shown in t11, TERM will fire top-down. It subsequently, through feedback, activates all meanings of the category type which predicted the TERM, all VTP type meanings in this case. This excitatory feedback, raises both VTP meanings in t12, for saw (d), as well as, building (e). However, the activity level of "building does" not reach threshold because of previous disambiguation of its NOUN meaning. When the VTP meaning, saw, fires (d) in t12, additional composition occurs. The VTP interpretation composes with a suitable TERM (a), one which matches feature attribute specifications of saw,




Figure 2.




Figure 3.

to produce a VIP type at t13 this will subsequently produce feedback at t14. Neither are shown.

3.5.3 Effect of a Different Grammar State on Processing

The final example, Figure 3, will use one of the "lesion" simulations using HOPE. The grammar representations remain intact. This example will present the understanding of the first three words of the sentence under the condition that they are presented faster than the system is processing. Effectively, a slow-down of activation spread to the grammar is assumed. Figures such as Figure 1 and Figure 3 can be compared to suggest possible language performance problems and to gain insights into their possible causes.

In Figure 3, when /TH-UH/ is introduced at t1 (a), all meanings are activated (b) as in Figure 1. The spread of activation to the grammar occurs in t2 (c). However, the second word, /B-OY/ (d) is "heard" at the same time as the activity reaches the grammar. The predictive activation spread from the grammar takes effect at t3, when the new word /S-AO/ (e) is "heard." The immediate result is that the NOUN meaning, saw (f), fires and is interpreted at t4 (g).

This shows in a very simple case, how the grammar can affect the processing states of an interactive parallel model. Timing can be seen to be critical. There are many more critical results that occur in such "lesion" simulations that better illustrate such grammatical affects, however they are very difficult to present in a static form, other than within a behavioral analysis of the overall linguistic performance of the entire model. This is considered an hypothesized patient profile and is described in Gigley (1985). Other examples of processing are presented in detail in Gigley (1982b; 1983).

3.6 Summary

The above figures present a very simple example of the interactive process. It is hoped that they provide an idea of the interactions and feedback, feedforward processing that is coordinated by the state of the grammar. Any prediction in the grammar that is not sufficiently active affects the process. Any decay that accidentally reduces a grammatical aspect can affect the process. The timing of activation, the categorical content and the interactions between interpretation and prediction are important factors when one considers grammar as part of a functioning dynamic system.

Finally, the Categorical Grammar is one form of a Context-Free (CF) grammar which provides a suitable integration of syntactic and semantic processing. In addition, it has been used in many studies of English so that instances of grammars sufficiently defined for the current implementation level of processing could be found. Other forms of grammar, such as Lexical-Functional Grammar (Kaplan and Bresnan, 1982) or Generalized

Phrase Structure Grammar (Gazdar, 1982; 1983) could be equally suitable.

The criteria to be met all that they can be encoded as predictive mechanisms, not necessarily unambiguous or deterministic, and also that they specify constraints on compositionality. The composition depends on adequate definition of interpretation constraints to assure that it is "computed" properly or else suitably marked for its deviation.

4. Conclusion

HOPE provides evidence for how one can view a grammar as an integrated part of a neurally-motivated processing model that is psychologically valid. Suitable constraints on grammatical form that are relevant for using any grammar in the CN context are that the grammar make serial predictions and provide the synchronization information to coordinate top-down effects of interpretation with the on-going process.

This type of model suggests that universals of language are inseparable from how they are computed. Universals of language may only be definable within neural substrata and their processes. Furthermore, if this view of linguistic universals holds, then grammar becomes a control representation that synchronizes the kinds of signals that occur and when they get propagated. The states of the grammar in this suggested view of grammatical function are a form of the rewrite rules that are the focus of much linguistic theory.

A neurally motivated processing paradigm for natural language processing, demonstrates how one can view an integrated process for language that employs integrated syntactic and semantic processing which relies on a suitable grammatical form that coordinates the processes.

5. Acknowledgements

The initial development of the reported research was supported by an Alfred P. Sloan Foundation Grant for "A Training Program in Cognitive Science" at the University of Massachusetts at Amherst. Continuing development is supported through a Biomedical Research Support Grant at the University of New Hampshire.

6. References

- Ajdukiewicz, D. Die Syntaktische Konnexität, 1935. Translated as "Syntactic Connection" in Polish Logic, S. McCall, Oxford, 1967, 207-231.
- Arbib, M.A. and Caplan, D. Neurolinguistics Must Be Computational. Behavioral and Brain Sciences, 1979, 2, 449-483.
- Collins, A.M., and Loftus, E.A. A spreading activation theory of semantic processing. Psychological Review, 1975, 82:6, 407-428.

Cottrell, G.W. and Small, S.L. A Connectionist Scheme for Modelling Word Sense Disambiguation. Cognition and Brain Theory, 1983, 6:1, 89-120.

Cottrell, G.W. A model of Lexical Access of Ambiguous Words. Proceedings of AAAI -- 1984.

Gazdar, G. Phrase Structure Grammar. In P. Jacobson and G. Pullum (eds.), The Nature of Syntactic Representation. Reidel, Dordrecht, 1982.

Gazdar, G. Phrase Structure Grammars and Natural Languages. Proceedings of the Eighth International Joint Conference on Artificial Intelligence. Karlsruhe, West Germany, 1983.

Gigley, H.M. Neurolinguistically Based Modeling of Natural Language Processing. Paper presented at the ACL Session of the Linguistic Society of American Meeting, New York, December, 1981.

Gigley, H.M. A computational neurolinguistic approach to processing models of sentence comprehension. COINS Technical Report 82-9, University of Massachusetts, Amherst, 1982.

Gigley, H.M. Neurolinguistically constrained simulation of sentence comprehension: Integrating artificial intelligence and brain theory. Ph.D. Dissertation, University of Massachusetts, Amherst, 1982b.

Gigley, H.M. HOPE -- AI and the Dynamic Process of Language Behavior. Cognition and Brain Theory, 1983, 6, 1.

Gigley, H.M. Computational Neurolinguistics - What is it all about? Proceedings of IJCAI 85, Los Angeles, to appear.

Gigley, H.M. From HOPE en l'ESPERANCE -- On the Role of Computational Neurolinguistics in Cross-Language Studies. Proceedings of COLING 84. Stanford University, July, 1984.

Goodenough, C., Zurif, E. and Weintraub, S. Aphasic's attention to grammatical morphemes. Language and Speech, 1977, 11-19.

Goodglass, H. Agrammatism. In H. Whitaker and H.A. Whitaker (eds.) Studies in Neurolinguistics, Vol. 1, Academic Press, 1976, 237-260.

Goodglass, H. and Berko, J. Agrammatism and inflectional morphology in English. Journal of Speech and Hearing Research, 1960, 3, 257-267.

Goodglass, H., Gleason, J., Bernholtz, N. and Hyde, M. Some Linguistic Structures in the Speech of a Broca's Aphasic. Cortex, 1970, 8, 191-212.

Hinton, G.E. Implementing Semantic Nets in Parallel Hardware. In G.E. Hinton and J.A. Anderson (eds.), Parallel Models of Associative Memory. Lawrence Erlbaum Associates Publishers, 1981.

Kaplan, R.M. and Bresnan, J. Lexical-Functional Grammar: A Formal System for Grammatical Representation. In J. Bresnan (ed.), The Mental Representation of Grammatical Relations. MIT Press, 1982.

Lavorel, P.M. and Gigley, H.M. Elements pour une theorie generale des machines intelligentes. Intellectica, 1983, 7, 20-38.

Lewis, D. General Semantics. In Davidson and Harmon (eds.), Semantics of Natural Language, 1972, 169-218.

Quillian, M.R. Semantic Memory. In M. Minsky (ed.), Semantic Information Processing. Cambridge, Ma.: MIT Press, 1980.

Small, S., Cottrell, G., and Shastri, L. Toward connectionist parsing. Proceedings of the National Conference on Artificial Intelligence, Pittsburgh, PA: 1982, 18-20.

Waltz, D. and Pollack, J. Massively Parallel Parsing: A Strongly Interactive Model of Natural Language Interpretation. Cognitive Science. In press.

Zurif, E.B. and Blumstein, S.E. Language and the Brain. In M. Halle, J. Bresnan, and G. Miller, (eds.), Linguistic Theory and Psychological Reality. MIT Press, 1978.