C E N T E R F O R R E S E A R C H I N L A N G U A G E

June 1990 **Vol. 4, No. 3**

The newsletter of the Center for Research in Language, University of California, San Diego, La Jolla CA 92093. *(619) 534-2536*; electronic mail: *crl@amos.ucsd.edu*

• • •

CONTENTS

EDITOR'S NOTE

This newsletter is produced and distributed by the **CENTER FOR RESEARCH IN LANGUAGE**, a research center at the University of California, San Diego that unites the efforts of fields such as Cognitive Science, Linguistics, Psychology, Computer Science, Sociology, and Philosophy, all who share an interest in language. We feature papers related to language and cognition (1-10 pages, sent via email) and welcome response from friends and colleagues at UCSD as well as other institutions. Please forward correspondence to:

Teenie Matlock, Editor
Center for Research in Language, C-008
University of California, San Diego 92093
Telephone: (619) 534-2536
Email: *crl@amos.ucsd.edu*

SUBSCRIPTION INFORMATION

If you are currently receiving a hardcopy of the newsletter only and have an email address to which the newsletter can be sent, please forward that information to CRL.

Can we send you the email version only to save printing and mailing costs? If you require a hardcopy in addition, please request it and we will be happy to send you one.

If you require the unformatted nroff version of the newsletter, please request it from CRL after you have received the current regular formatted version.

If you know of others who would be interested in receiving the newsletter, please forward the email or postal mailing address to CRL. Thank you.

BACK ISSUES

Back issues of this newsletter are available from **CRL** in hard copy as well as soft copy form. Papers featured in previous issues include the following:

The Cognitive Perspective
Ronald W. Langacker
Department of Linguistics, UCSD
vol. 1, no. 3, February 1987

Toward Connectionist Semantics
Garrison Cottrell
Institute for Cognitive Science, UCSD
vol. 1, no. 4, May 1987

Dimensions of Ambiguity
Peter Norvig
Computer Science, UC Berkeley
vol. 1, no. 6, July 1987

Where is Chomsky's Bottleneck?
S.-Y. Kuroda
Department of Linguistics, UCSD
vol. 1, no. 7, September 1987

Transitivity and the Lexicon
Sally Rice
Department of Linguistics, UCSD
vol. 2, no. 2, December 1987

Formal Semantics, Pragmatics, and Situated Meaning
Aaron Cicourel
Department of Sociology, UCSD
vol. 2, no. 3, January 1988

Rules and Regularities in the Acquisition of the English Past Tense
Virginia Marchman
Department of Psychology, UC Berkeley
vol. 2, no. 4, April 1988

A Geometric Conception of Grammar
S.-Y. Kuroda
Department of Linguistics, UCSD
vol. 2, no. 5, June 1988

Harris and the Reality of Language
S.-Y. Kuroda
Department of Linguistics, UCSD
vol. 3, no. 1, September 1988

A Connectionist Perspective on Prosodic Structure
Mary Hare, Dept. of Linguistics, UCSD
David Corina, Dept. of Psychology, UCSD
G.W. Cottrell, Dept. of Computer Science, UCSD
vol. 3, no. 2, November, 1988

Competent Scientist Meets the Empiricist Mind
Valerie Walker
Department of Philosophy, UCSD
vol. 3, no. 4, April, 1988

Foreign Policy By Metaphor
Paul Chilton, University of Warwick, U.K.
George Lakoff, UC Berkeley
vol. 3, no. 5, June, 1989

Guarani Possessive Constructions
Maura Velazquez
Department of Linguistics, UCSD
vol. 3, no. 6, July, 1989

Anticipatory Coarticulation and Aphasia: Implications for Connectionist Models of Speech Production
William F. Katz
Department of Psychiatry, UCSD
vol. 4, no. 1, December, 1989

Toward a Connectionist Representation of Grammatical Knowledge
Catherine L. Harris
Department of Cognitive Science, UCSD
vol. 4, no. 2, January, 1990

Summer Events at UCSD

The 1990 Connectionist Models Summer School is taking place at UCSD this summer, from June 22 to July 1. This is the third in a series which has been held every two years, starting in 1986. The first two schools were held at Carnegie Mellon University; the third and the fourth (in 1992) will be held at UCSD.

The summer school brings together 50 graduate students and 30 faculty from around the world for a series of lectures, workshops, and informal discussions about connectionist models in a wide range of domains, including cognitive science, neuroscience, computer science, linguistics, philosophy, biology, engineering, mathematics, and physics. The summer school is sponsored by the Institute for Neural Computation, with cooperation from the Center in Research in Language and the Department of Cognitive Science.

Proceedings of the summer school, edited by D. Touretzky, J. Elman, T. Sejnowski, & G. Hinton, will be published by Morgan Kauffman in the fall.

The Rule Controversy: A Cognitive Grammar Perspective

Ronald Langacker
University of California, San Diego

Few issues are more consequential for our view of language and mind than the nature and status of linguistic rules. The connectionist challenge to the classic generative conception of rules should therefore be welcomed by all concerned, whether because it ultimately shows that explicit rules are superfluous, or because it stimulates generative theorists to refine their notion of rules and the basis for postulating them. My own stand on this issue represents a kind of middle ground. While cognitive grammar has natural affinities to connectionism, it does posit something comparable to rules. At the same time, it conceives of these entities in a way that makes them amenable to an interpretation in terms of connectionist processing. My objectives here are thus to sketch this conception of linguistic rules, to portray it as the logical culmination of trends observable within the generative tradition itself, and thereby to point the way to a possible eventual convergence.

1. Basic Features of Cognitive Grammar

While I can offer here only the briefest of characterizations, I maintain that the theory of **cognitive grammar** is intrinsically desirable on grounds of naturalness, conceptual unification, and theoretical austerity.[1] One aspect of its naturalness resides in the fact that it posits only **semantic structures**, **phonological structures**, and **symbolic links** between the two. This basic organizational feature correlates directly with the primary function of language, that of permitting meanings to be symbolized by phonological sequences. The theory achieves conceptual unification by claiming that grammar is fully reducible to symbolic relationships: lexicon, morphology, and syntax form a continuum exhaustively describable in terms of **symbolic structures**, each comprising both a semantic and a phonological structure together with the symbolic relationship that links them. Grammar is therefore said to be **symbolic** in the specific sense that it **reduces to form-meaning pairings**.[2]

Cognitive grammar is theoretically austere by virtue of the **content requirement**, which stipulates that only three kinds of elements are ascribable to a linguistic system: (i) semantic, phonological, and symbolic structures that occur overtly as (parts of) expressions; (ii) schematizations of permitted structures; and (iii) categorizing relationships between permitted structures. To take a phonological example, *sim*, *lep*, and *tich* are syllables that occur overtly in English; [CVC] is a schematization over such structures; and [[CVC]===>[sim]] represents the categorization of [sim] as an instance of [CVC]. Similarly, the words *head*, *heart*, *hand*, *hip*, and *heel* are overtly occurring symbolic structures; the schema [BODY PART/hV(C)C] expresses their commonality and thereby defines a class of symbolic elements; and [[BODY PART/hV(C)C]===>[HEAD/head]] indicates that *head* is a member of that class.[3] The effect of the content requirement is to rule out any descriptive construct not straightforwardly derivable, by the well-attested processes of schematization and categorization, from structures experienced directly (i.e. meanings and phonological sequences).

In this framework, therefore, linguistic regularities take the form of **schemas**. What do I mean by that term? Importantly, a schema is not the same as a list of criterial attributes or a bundle of separate features. It is rather a coherent, integrated structure comparable in most respects to those which support its extraction. A schema's internal organization is precisely parallel to that of the semantic, phonological, or symbolic structures it schematizes, thus reflecting whatever commonality they exhibit. It does however abstract away from their points of divergence, being neutral or less specific in regard to each; overall, then, it is characterized at a lower degree of precision and detail. The essential component of this relationship is **granularity**: relative to the fine-grained specifications of the structures it categorizes, a schema provides only a coarse-grained characterization in which certain detailed features fail to appear owing to insufficient resolution.[4] A difference in levels of schematicity (or specificity) is the basis for **hyponymy** in lexical semantics, and more generally, for the relation between **superordinate** and **subordinate** levels in a taxonomic hierarchy.

I will indicate a relationship of schematicity by means of a double arrow. Thus, *A*===>*B* states that *A* is schematic for *B*, or conversely, that *B* **elaborates** (or **instantiates**) *A*. Another basic type of categorization, indicated by a single arrow, is **extension** from a **prototype**. Hence *A*--->*B*

states that *A* categorizes *B* despite some conflict in specifications. Linguistic categories are typically **complex**, in the sense that their proper characterization takes the form of a **network** whose nodes are structures linked by categorizing relationships of these two sorts. The structures that participate in these relations (the nodes in the network) can be semantic, phonological, or symbolic, and may exhibit any degree of internal complexity. The network may, for example, comprise alternate senses of a lexical item, or the various allophones of a phoneme. The individual nodes of the network may even consist of entire grammatical constructions, analyzed here as configurations of symbolic structures (cf. Langacker 1988c; Lakoff 1987, case study 3).

The possibility of reducing grammar to configurations of symbolic structures presupposes an appropriate view of linguistic semantics, one that equates meaning with **conceptualization** and properly accommodates **construal**, i.e. our capacity for conceiving the same situation in alternate ways. One aspect of construal is our ability to conceptualize an entity at various levels of specificity, as reflected in **schematic hierarchies** such as [THING] ===> [ANIMAL] ===> [MAMMAL] ===> [DOG] ===> [BEAGLE] or [SEGMENT] ===> [CONSONANT] ===> [STOP] ===> [VELAR STOP] ===> [k]. We are further able to conceptualize a situation from different perspectives (consider *come* vs. *go*), to construe it in relation to different background assumptions and expectations (e.g. *half-empty* vs. *half-full*), and to render certain entities more prominent than others. One type of prominence is figure/ground organization. Normally a relational expression accords one of its participants the status of figure within the scene; I call this the **trajector**, and refer to an entity of secondary salience as a **landmark**. Thus the semantic contrast between *X is near Y* and *Y is near X* resides in whether *X* is construed as the trajector (relational figure) and is being located with reference to *Y*, or conversely. I use the term **profiling** for a second type of prominence, whereby every expression--within the conception it evokes--singles out a particular substructure as a kind of focal point; this substructure (the **profile**) is the entity that the expression **designates**. For instance, *knuckle* evokes the conception of a finger and profiles (designates) one of its joints. *Near* profiles the **relationship** between two conceived entities (its trajector and landmark).

Given a conceptualist semantics based on construal, the reduction of grammatical structure to symbolic relationships becomes quite feasible. Grammar can, I believe, be fully and revealingly characterized using only symbolic structures (form-meaning pairings); on this view grammar is indissociable from meaning, and all grammatical elements are attributed some kind of conceptual import. For instance, an expression's grammatical category is determined by the nature of its profile (Langacker 1987a, 1987b). Thus, an expression is categorized as a noun by virtue of profiling a **thing** (under a highly abstract definition of that term). Rather than being marked as such by a diacritic or syntactic feature (devices not permitted by the content requirement), a word like *cat* is inherently classed as a noun because it instantiates the schema defining the category: [[THING/...]===>[CAT/cat]].[5] Likewise, every verb is claimed to profile a **process** (defined as a relation scanned sequentially in its evolution through conceived time), e.g. [[PROCESS/...]===>[SPRAY/spray]]. Semantic characterizations have also been proposed for many other basic grammatical notions, including noun phrase, finite clause, head, complement, modifier, coordination, subordination, auxiliary verb, subject, object, transitivity, unaccusative, and ergativity (Langacker 1986b, 1986c, 1987a, 1989, 1990, *in press*, *to appear*; Rice 1987b).

Our focus here is on **grammatical rules**, which pertain to the combination of simpler symbolic structures to form more complex ones. I say that a **symbolically complex** expression is formed by the **integration** of two or more **component structures** to yield a **composite structure**. For example, the two component structures, *jar* and *lid*, can be integrated to form the composite structure *jar lid*. Each is symbolic, comprising a semantic structure symbolized by a phonological structure: [JAR/jar], [LID/lid], [JAR LID/jar lid]. Semantically, *jar* designates a particular type of container, while *lid* profiles the cover to a container characterized only schematically. Phonologically, *jar* and *lid* are each characterized as words. Integration is effected by **correspondences** established between subparts of the component semantic and phonological structures. Thus, the specific container profiled by *jar* is put in correspondence with the schematic container evoked by *lid*, and *lid* is identified as the word that directly follows *jar* in temporal order. The composite structure is formed by superimposing corresponding entities and merging their specifications. The composite structure [JAR LID/jar lid] inherits the profiling of *lid*, which is therefore the **head**.

Jar lid instantiates a general pattern of compound formation in English. The regularity exhibited by *jar lid*, *door knob*, *pencil eraser*, *milk*

carton, *sea captain*, *salad oil*, *school bus*, *garbage man*, *train station*, *delivery boy*, and countless others is captured by an appropriate schema, which expresses their commonality while abstracting away from their points of contrast. Since one aspect of their commonality lies in their symbolic complexity and the nature of the relationships holding among the component and composite structures, those organizational features are reflected in the schema. This **constructional schema** is directly analogous to the complex expressions from which it is extracted: the only difference is that the component and composite structures are themselves schematic rather than specific. Semantically, for instance, the first component is merely characterized as profiling a thing (this makes it a noun), and the second, as profiling a thing that bears some unspecified relationship to another thing, which remains unprofiled. However, the profile of the first component corresponds to the unprofiled thing within the second (just as in *jar lid*), and the profile of the second prevails at the composite-structure level. This latter feature represents the generalization that the second element of English noun-noun compounds functions as the head.

In principle, any regularity observable across a set of complex expressions can be captured by an appropriate constructional schema. Such schemas are cognitive grammar's equivalent of grammatical rules. As complex symbolic structures derived by schematization from overtly-occurring expressions, they clearly conform to the content requirement. Moreover, they represent established patterns potentially available for the categorization of novel expressions; they can also be viewed as templates employed in the assembly or evaluation of such expressions. A particular expression is simultaneously categorized by numerous schemas, each pertaining to a different aspect of its structure. The full set of categorizing relationships in which an expression thusly participates constitutes its **structural description** and determines its status vis-a-vis the linguistic system. An expression is **conventional** (i.e. well-formed) with respect to a given aspect of its structure just in case it elaborates the categorizing schema (i.e. *A*===>*B*) rather than conflicting with its specifications (*A*--->*B*).

Hence the entire complex structure represented by *jar lid* is categorized by the constructional schema, and to the extent that *jar lid* constitutes a fixed, familiar expression, both it and the categorizing relationship belong to the grammar of English (i.e. they are facets of a speaker's grasp of established convention). These elements are part of a substantial network describing the various patterns and subpatterns of English compounding. The nodes of this network--linked by categorizing relationships of elaboration or extension--include specific compounds with the status of fixed expressions, subschemas characterized at different levels of specificity, and higher-level schemas representing broad generalizations. These structures have varying degrees of cognitive salience, in the sense of "entrenchment" or ease of activation; thus accommodated is the distinction between "major" and "minor" patterns, or productive vs. non-productive rules. It can be seen, therefore, that a network of this sort both captures any discernible regularities and specifies how general patterns are actually implemented in conventional usage. Its configuration (including salience) determines which particular instantiations--out of the vast range that a higher-level schema would theoretically permit--actually tend to be used, and with what degree of likelihood. Distributional restrictions are thus handled in a cognitively plausible way, without resorting to devices that would violate the content requirement (e.g. rule features or diacritics).

Though I can hardly prove it here, I believe this view of linguistic structure to be both viable and revelatory, and will assume its workability for purposes of the following discussion. Its basic affinity with connectionism should be readily apparent. First, it reduces grammar to form-meaning pairings, which PDP models are well suited to deal with. Second, it makes no qualitative distinction between rules and data; schemas and their instantiations differ only in level of specificity, which is a matter of degree. Third, it relies on the extraction of generalizations from positive instances, through the reinforcement of common organizational features. Finally, it considers local regularities to be at least as significant in language processing as high-level generalizations. Having noted these similarities, let us now consider how cognitive grammar's conception of linguistic rules relates to developments within the generative tradition.

2. A Spectrum of Positions

To put the issue in proper perspective, let us outline two extreme positions concerning the cognitive representation of linguistic structure. These positions are caricatures--no serious scholar subscribes to either. They do however define the endpoints in a spectrum of possible views and are therefore useful as reference points. We can think

of them as corresponding to the most simplistic notion that an uninformed generativist or connectionist might conceivably entertain concerning what the other believes. In this spirit of misconception and caricature, we can imagine a connectionist referring to the position he wrongly imputes to a generativist as "empty symbol pushing". We can similarly imagine a generativist offering for his misinterpretation of connectionism the catchy descriptive label "mind as mush".

"Empty symbol pushing" is the (hypothetical) position that a language is fully describable by a set of rewriting rules very much like the grammars of simple artificial languages found in Chomsky's early writings (e.g. 1957, 1965; Chomsky and Miller 1963). Except for a vast difference in complexity, a natural-language grammar is thought of as being precisely analogous, say, to that comprising the two rewriting rules *S*--->*aSb* and *S*--->*ab*, which generates all and only the sentences consisting of a string of *a*'s followed by a string of *b*'s of equal length. Grammars of this kind have four essential properties: (i) They are **explicit** in the strong sense of being **generative**, i.e. they offer a full and precise mathematical characterization of all and only the well-formed sentences of a language. (ii) They are **constructive**, consisting of rules for assembling grammatical expressions. Observe that these rules are formally distinct from the expressions they generate (e.g. *S*--->*aSb* is formally dissimilar to *aaaabbbb*). (iii) Every rule is **fully general**; it applies in the derivation of an open-ended set of sentences, and is applicable without exception to any structure that meets its specifications. (iv) The rules manipulate **contentless symbols**. A non-terminal symbol (such as *S*) has neither semantic nor phonological content. And while terminal symbols (such as *a* and *b*) might be thought of as vocabulary items, their meaning is irrelevant to grammatical derivation--a string's grammaticality can be determined from its form alone (e.g. by counting the number of *a*'s and *b*'s).

At the other extreme, "mind as mush" describes the (hypothetical) view that mental processing has no interesting structure whatever--it merely involves an array of faceless units behaving in a squishy fashion. With sufficient training, a PDP system comes to function in a way that mirrors the statistical regularities inherent in its input, and that is all the investigator is concerned with accomplishing. Since learning is confined to adjustments in connection weights, and since no individual weight or adjustment can be identified with a specific linguistic structure, there is no point in searching for anything that might correspond to the discrete constructs posited by linguists. In particular, the explicit rules they formulate are considered superfluous and fictitious--after all, a PDP system accommodates linguistic regularities without resorting to such entities.

Thus the "empty symbol pushing" and "mind as mush" positions represent polar opposites that are clearly irreconcilable. They do not represent positions actually espoused, however, and as we move away from these caricatures to more accurate descriptions of what generativists and connectionists currently believe, the contrast appears considerably less stark.

Now as I understand the connectionist enterprise, its objective is a realistic model of cognitive processing that accurately reflects observed behavior displaying all degrees of regularity, from the idiosyncratic to the exceptionless. Importantly, PDP systems are quite capable of structured behavior which is "rule-governed" at least in the sense that specific patterns of activation are crisply and reliably elicited by particular kinds of input. Connectionists are very much concerned, moreover, with finding out just how their systems work. It is for this reason that they monitor the activation of hidden units to see what features they serve to detect; perform cluster analyses to determine, on the basis of response similarity, the implicit categorization the system has imposed on the input data; consider its behavior in terms of locations and trajectories in state space; and so on. The functioning of a connectionist model is therefore regarded as organized activity susceptible to being studied and understood, and it is quite conceivable that certain aspects of this processing might be identified with particular linguistic constructs.

Likewise, contemporary generative theory bears little resemblance to the "empty symbol pushing" caricature. With respect to all four properties noted previously, trends within the generative tradition have moved it closer to an outlook having a certain amount of commonality with the connectionist perspective. (i) Although explicit rules and representations are still considered *de rigueur*, no longer is a grammar universally conceived as a strictly generative, algorithmic device. In fact, theories (notably government-binding theory) are now proposed and developed without any serious attempt at formalization. (ii) It is increasingly less common for theorists to use rewriting rules or to conceive of grammars as constructive devices. There is more emphasis on surface constructions (as opposed to derivations from underlying structures), and characterizations involving the simultaneous satisfaction of multiple

constraints are envisaged both in "unification-based" approaches and in discussions of "modularity". (iii) That a linguistic description requires statements at all levels of generality is fully recognized, and a variety of formal devices have been proposed for exceptions, irregularity, and rules of limited productivity. These are often consigned to "the lexicon", which is now regarded as being of prime importance. (iv) The intimate association of grammar and meaning is coming to be appreciated. It has long been acknowledged that grammaticality judgments cannot be based on strings of words per se, but pertain to particular structural descriptions (if not specific interpretations). There is also a growing realization that differences in grammatical behavior correlate with differences in meaning, and that semantic (or at least "pragmatic") consequences follow from the choice among alternative constructions.

Despite these developments, the distance between the generativist and connectionist outlooks remains substantial, and differences in underlying philosophy may well preclude an eventual convergence.[6] This is where cognitive grammar enters the picture, for in a sense it provides a bridge between the two. Certain affinities to connectionism have already been pointed out. Now I will not assert that cognitive grammar shows comparable affinities with generative grammar; some radical adjustments in generative thought would be needed to bring them into alignment. Nevertheless, cognitive grammar is not unreasonably viewed as representing the logical culmination of all the aforementioned trends, the kind of natural, unified, and restrictive theory that might emerge if they ran their full course and certain basic but erroneous assumptions were abandoned.[7]

Cognitive grammar's position on these issues can be summarized as follows: (i) A grammar is specifically *not* conceived as a generative device. Since meaning is open-ended ("encyclopedic"--cf. Haiman 1980) and based on construal, one cannot envisage the recursive enumeration of "all and only the well-formed sentences [form-meaning pairings] of a language", for this is not a well-defined set. Moreover, there is no expectation that any single formalism or representational format will prove uniquely appropriate for describing a particular aspect of linguistic structure, or capable of providing an exhaustive account.[8] (ii) The grammar of a language is not thought of as *constructing* expressions (giving them as "output"), but simply as an inventory of conventional structures available for their *categorization*. An expression's structural description resides in simultaneous categorization by numerous schemas, each amounting to a constraint pertaining to some aspect of its organization. Hence cognitive grammar is a "unification-based" model *par excellence*. (iii) It is also a **usage-based model** (Langacker 1988c), by which I mean that considerable emphasis is placed on specific expressions and low-level generalizations. A speaker's linguistic knowledge subsumes a vast set of fixed expressions--not just lexical items in the usual sense, but also standard collocations, formulaic expressions, and all manner of complex locutions representing the normal way of phrasing things in the language. Rules are merely schematizations of expressions; they represent all levels of generality, and coexist in the grammar with any instantiating expressions that are learned and familiar. Moreover, since schemas compete for activation (i.e. for the privilege of categorization and structural description) on the basis of specificity as well as entrenchment, lower-level schemas are essential to linguistic structure, serving as the primary locus of distributional information. (iv) In this framework, grammatical structure reduces to the structuring and symbolization of conceptual content. Meaning and grammar are not just intimately associated, therefore, but *inherently indissociable*.

To generative theorists the non-generative and non-constructive nature of cognitive grammar should no longer seem exotic or unduly bothersome, and its greater emphasis on low-level generalizations is primarily a matter of degree. What about the idea that rules are just schematized expressions? Though more likely to be resisted, this too is based on notions also encountered in the generative tradition (structural templates; multiple constraint satisfaction); adopting it would simply be a matter of recognizing their universal applicability. It is much harder to imagine generativists ever accepting the claim that grammar reduces to symbolic relationships, so fundamental to their world view is the autonomy thesis. I would argue, however, that consideration of the autonomy thesis has been clouded by a certain amount of conceptual unclarity and the confounding of distinct issues. When these are properly sorted out, it becomes possible to accommodate the valid observations that have been taken as sustaining that thesis, while at the same time achieving the reduction of grammar to configurations of symbolic structures. In sum, even though these basic properties of cognitive grammar may seem quite radical from the standpoint of generative theory, they have some precedent in that tradition and are not entirely unresponsive to its concerns. And it is these properties that offer the realistic prospect of

a connectionist interpretation.

3. The Nature of Rules

If the generativist position is that explicit rules are needed for the proper characterization of linguistic structure, while the connectionist position is that they are not, I can summarize my own view by saying that I agree with them both. The apparent contradiction is resolvable once it is realized that the generative and PDP programs stem from different initial concerns, and that certain statements made in regard to rules pertain to distinct issues and are therefore incommensurate.

Despite its concern with psychological questions, we should not forget that generative grammar grew out of the tradition of **descriptive linguistics**, whose goal was the recording and analysis of a language, followed by an *explicit description* of its structure for the benefit of other investigators. At least in the early days, writing a (partial) transformational grammar of language X was considered a reasonable goal, and for some languages such a grammar was the only substantial source of information. There was nothing inherently odd about the notion of dispatching a transformational grammarian to the field for purposes of describing an otherwise unknown language (though fieldwork tended not to be a primary emphasis). By contrast, faced with the imperative of recording and describing an unknown language on the verge of extinction, one would hardly think of sending a connectionist.

The task of describing linguistic structure in a usable form, for purposes of documentation and further analysis, thus provided the original context in which the generative notion of explicit rules was conceived and developed. This basic descriptive goal has no counterpart in the connectionist program, which is solely concerned with the nature of cognitive processing. Now if the converse were also true--that is, if generativists were solely concerned with description, and not at all with cognition--there would be no grounds for conflict. But of course there are, since one of Chomsky's major innovations was the proposal that linguistic descriptions be considered hypotheses about certain aspects of cognitive structure. With purposeful ambiguity, he used the term **grammar** for the cognitive representation of linguistic structure, as well for the linguist's attempt to describe it. Our interest here, though, is in sorting out the issues to see just where the conflict lies. I will therefore distinguish between an **internal** and an **external grammar**, i.e. between the mental representation of language (whatever its nature) and what linguists produce by way of its characterization.

The descriptive legacy is, I believe, one factor in the generative commitment to explicit rules. An *external* grammar has to be reasonably explicit in order to fulfill its purpose, whether this be practical or intellectual. In particular, explicit statements of linguistic regularities provide a characterization of what it is that a processing model has to account for (irrespective of whether it incorporates any direct analogs of those statements). One can of course argue about what form a description ought to take--how formal it needs to be, what kinds of constructs should be posited, how much regularity the data actually exhibits, etc. But if we confine our attention to external grammars serving a descriptive function, the validity and even the necessity of formulating explicit rules (of one sort or another) seems readily apparent. The controversial issue is whether such rules should also be ascribed to the *internal* grammar.

A second factor in the generative commitment to explicit rules is the autonomy thesis, the claim that grammar (or syntax in particular) constitutes a distinct level or component of linguistic structure with its own representations, primitives, and organizational principles. To sustain the autonomy thesis, it is argued that grammatical patterns and restrictions cannot be derived as automatic consequences of meaning or other independent factors--consequently they have to be *stated explicitly* as part of a linguistic description, and specifically learned in language acquisition (cf. Newmeyer 1983).

Let me first point out that this argument fails to establish the autonomy thesis, for it harbors a fallacy. The basic observation is certainly correct: while grammar can usually be seen as *motivated* on grounds of meaning or function, its specific detail is not in general *predictable* and must therefore be described by linguists and learned by speakers. But from this observation one cannot legitimately draw the further conclusion that grammar (or syntax) constitutes a distinct and autonomous component of the linguistic system. This further step confuses two issues that are in principle quite distinct, namely the *kinds of structures* that must be posited and the *predictability* of their behavior. Cognitive grammar is thus coherent and consistent in accepting the non-predictability of grammatical structure while nevertheless denying its autonomy. It acknowledges that grammatical patterns and restrictions must indeed be learned and explicitly described, but claims that their

proper characterization requires only symbolic structures (networks of constructional schemas). Rather than being autonomous vis-a-vis semantics, grammar reduces to form-meaning pairings.

I thus consider the generative conception of explicit rules to be valid in certain respects but not all. A language does exhibit structural (including grammatical) regularities, many of which are "autonomous" in the limited sense that they do not follow inexorably as wholly predictable consequences of other factors. These regularities can reasonably be referred to as *rules*, and to serve its descriptive function an external grammar must state these rules explicitly. Moreover, since a speaker has to learn the patterns and restrictions of his language, comparable information must somehow be provided by the internal grammar, i.e. it must have some kind of cognitive instantiation. The form this knowledge takes, however, may be quite different from what generativists tend to assume; my central point is that cognitive grammar affords a new and useful perspective on this question. In particular, it offers a distinct conception of linguistic rules that is not inherently incompatible with either the letter or the spirit of connectionism.

Cognitive grammar recognizes the importance of explicit description, both for practical reasons and as an essential step toward determining the mental representation of linguistic structure. Various formats have been adopted for descriptive purposes, each revelatory in its own way, but no single format is considered uniquely privileged or presumed capable of capturing every significant aspect of a phenomenon (cf. fn. 8). There is also no supposition that any particular notation or descriptive device translates directly into claims about the basic nature of cognitive processing.[9] Descriptions are attributed the more limited role of elucidating certain regularities that we can reasonably expect to be reflected (and hopefully discernible) somewhere within the mental processing constitutive of linguistic ability. To the extent that processing regularities correspond to the kinds of patterns linguists seek to discover and describe, they can be regarded as the cognitive embodiment of linguistic rules. Two basic questions then arise: What in fact is their nature? And where can we expect to find them?

Their nature is indicated by two fundamental claims of the theory: that grammar reduces to configurations of symbolic structures (form-meaning pairings); and that rules are merely schematizations of expressions (coarse-grained characterizations representing the commonality that emerges at an appropriate level of abstraction). The first claim implies that the rules of the internal grammar neither comprise nor manipulate contentless symbols. All linguistic structures are either semantic, phonological, or symbolic. Being symbolic, grammatical structures have both semantic and phonological value (though it may be quite abstract). The second claim entails the absence of any qualitative distinction between rules and expressions--apart from their level of specificity, generalizations have the same basic character as the data they account for. Now if rules are conceived in this fashion, they should pose no problems for a PDP system (provided that it is capable of representing both semantic and phonological structures). The generalizations extracted by such a system can, I suggest, be identified with the schemas posited in cognitive grammar.

Where can these rules (schemas) be found? How can their postulation be reconciled with the connectionist assertion--waved like a red cape in front of generativists--that linguistic regularities can be handled by PDP systems which make no use of explicit rules? We must first be clear about what is actually intended by this assertion. Its import is twofold: that a system's specific activity is not directed by a central program (a list of statements telling each unit what to do at each step); and that information is stored exclusively in connection weights (none of which can be equated with any particular linguistic construct). Now it would seem that this abnegation of rules precludes their ascription to the internal grammar. There is indeed an incompatibility if one insists that rules are directly analogous to the instructions of a computer program, or that rules are stored as such at the most basic level of representation. Rules can, however, be accorded a very different status, in which case they do not run afoul of the connectionist prohibition but are simply incommensurate with it. In cognitive grammar, rules are conceived as *regularities in the mental processing constitutive of linguistic ability*. They are consequently *emergent* rather than fundamental; instead of being separately stored or represented in the form of instructions, they are inherent in the system's processing activity. Interpreted as recurrent patterns of neural activation, rules are wholly consistent with connectionist principles.

Let us consider this conception of rules and their PDP implementation in somewhat more detail. At the most basic level, linguistic knowledge is stored in connection weights. Neither rules nor any other linguistic elements are directly or individually discernible at that level, however. To find the cognitive correlates of

linguistic constructs, we must instead look at higher levels of organization, and specifically at the patterns of activation constrained by those weights. The evocation of a particular linguistic structure--be it semantic, phonological, or symbolic--resides in the occurrence of a particular pattern of neural activation. This pattern may be relatively simple, or it may be extraordinarily complex, comprising elaborately architectured cascades of activity involving many populations of units over a substantial span of processing time.[10] But regardless of complexity, a pattern is describable as either a **location** in the **state space** defined by the activation levels of the system's units, or else a **trajectory** through state space (i.e. a series of locations).

The structures that concern us are schemas and their instantiating expressions. It is crucial that the notation employed for their relationship, namely *A===>B*, not be construed as making a specific claim about the nature of its cognitive representation. The practice of using distinct symbols for a schema and its instantiation is helpful (even necessary) for analytical and descriptive purposes, but it is not meant to imply that they are discrete and separate psychological entities. Rather, I conceive of a schema as being **immanent** in its instantiations, i.e. as inherent in (and shared by) the activation patterns in which its instantiations reside. What does it mean, exactly, for one structure (or activation pattern) to be immanent in another? One way to interpret it in connectionist terms pertains to locations and trajectories in state space.[11] A location can be characterized with varying degrees of precision, being point-like or diffuse depending on whether activation levels are specified quite narrowly or only as falling within certain bands of values. And as a series of locations, a trajectory can similarly be characterized with varying degrees of exactitude (as either line-like or swath-like). We can thus describe the relationship between a schema and its instantiations as one of *inclusion in state space*: a schema corresponds to a diffuse region (or swath-like trajectory) in state space, and each instantiation to a point-like region (or line-like trajectory) contained within it.

On this account, the extraction of schemas is a non-mysterious process which results in essentially automatic fashion from the use of instantiating expressions. An expression's occurrence tends to strengthen the connection weights responsible for the pattern of activation that it comprises, and thus to facilitate the subsequent occurrence of another pattern in the same general region of state space. Hence the frequent use of expressions sufficiently similar that they cluster in such a region induces an adjustment of the responsible weights which renders the occurrence of any pattern falling within that region more likely or easily elicited than it would otherwise be. That adjustment constitutes the extraction of a schema. A schema is immanent in its instantiations in the sense that being located in a point-like region of state space entails being located in a diffuse region that encompasses it. Moreover, by virtue of facilitating a pattern falling anywhere within that region, a schema has an active, causal role in cognitive processing--it is not epiphenomenal, unless one wishes to say that all linguistic entities are epiphenomenal.

Let me conclude by noting certain challenges that cognitive grammar poses for connectionist modeling. First, to be linguistically viable a PDP system must be capable of representing *structured conceptualizations* of extraordinary intricacy (see especially Langacker *to appear*). Second, it must allow distinct structures to be co-activated and linked by correspondences while to some degree retaining their separate identity. This is needed for the characterization of grammatical constructions (cf. Figs. 1 and 2), metaphorical structuring (Lakoff and Johnson 1980), and correspondences between elements of different mental spaces (Fauconnier 1985). Finally, a linguistically adequate PDP system will have to accommodate the many dimensions of *construal* (Langacker 1988b), including such factors as profiling, figure-ground organization, and vantage point. While I have little doubt that these are all susceptible to connectionist treatment, not much attention has thus far been accorded them. Serious consideration of these matters would greatly facilitate a mutually instructive interaction between connectionism and cognitive linguistics.

Footnotes

[1]I have been developing this theory since 1976. By now, it has been described in numerous works and applied to a broad variety of representative linguistic phenomena. See, for example, Casad 1982; Casad and Langacker 1985; Cook 1988; Hawkins 1984; Janda 1984, *to appear*; Langacker 1982, 1984, 1985, 1986a, 1986b, 1986c, 1987a, 1987b, 1987c, 1988a, 1990, *in press*, *to appear*; Lindner 1981, 1982; Maldonado 1988; Poteet 1987; Rice 1987a, 1987b, 1988; Rudzka-Ostyn 1988; Smith 1985, 1987; Tuggy 1980, 1981, 1986, 1989; Vandeloise 1984, 1985a, 1985b, 1986, 1987.

[2]Crucially, this sense must not be confused with that implied in speaking of **symbolic** (as opposed to **connectionist**) accounts of cognitive processing.

[3]This particular class happens not to be structurally significant, but a comparable class might very well be. In a given language, for instance, the class of body-part terms conforming to a certain phonological pattern might all form their plurals in the same way.

[4]Another way to put it is that a schema allows a wider range of values along some or all parameters of the characterization.

[5]That is, a noun is characterized schematically as an expression that designates a thing (defined abstractly) and is manifested phonologically as any kind of phonological sequence.

[6]I refer here to a true *rapprochement*, not such half-way measures as implementing a standard generative analysis in a PDP model, or using rule-based and connectionist accounts for different components (e.g. syntax vs. lexicon, or competence vs. performance).

[7]As a historical note, I should point out that the basic framework of cognitive grammar has been in place for well over a decade--it has for the most part *anticipated* rather than *followed* these trends.

[8]Though I would not want to push the metaphor too far, it is useful in this regard to think of language as being analogous to a biological organism--however thoroughly it might be described, further characterization (in finer detail or from another perspective) can still be both valid and revelatory. (Cf. Langacker *to appear*, 12.1.)

[9]Hence a pictorial representation does not imply that the brain stores information in the form of pictures. Similarly, a formulaic representation does not imply that cognitive processing involves formulas or the manipulation of discrete symbols.

[10]A pattern of either sort is referred to as a **cognitive event** in Langacker 1987a.

[11]This interpretation was suggested to me by Steve Poteet.

Bibliography

Casad, Eugene H. 1982. *Cora Locationals and Structured Imagery*. San Diego: UCSD doctoral dissertation.

-----, and Ronald W. Langacker. 1985. '"Inside" and "Outside" in Cora Grammar'. *IJAL* 51.247-281.

Chomsky, Noam. 1957. *Syntactic Structures*. The Hague: Mouton.

-----. 1965. 'Three Models for the Description of Language', in R. Duncan Luce *et al.* (eds.), *Readings in Mathematical Psychology*, vol. 2, p.105-124. New York: Wiley.

-----, and George A. Miller. 1963. 'Introduction to the Formal Analysis of Natural Languages', in R. Duncan Luce *et al.* (eds.), *Handbook of Mathematical Psychology*, vol. 2, p.269-321. New York and London: Wiley.

Cook, Kenneth W. 1988. *A Cognitive Analysis of Grammatical Relations, Case, and Transitivity in Samoan*. San Diego: UCSD doctoral dissertation.

Fauconnier, Gilles. 1985. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, Mass. and London: MIT Press/Bradford.

Haiman, John. 1980. 'Dictionaries and Encyclopedias'. *Lingua* 50.329-357.

Hawkins, Bruce W. 1984. *The Semantics of English Spatial Prepositions*. San Diego: UCSD doctoral dissertation.

Janda, Laura A. 1984. *A Semantic Analysis of the Russian Verbal Prefixes ZA-, PERE-, DO-, and OT-*. Los Angeles: UCLA doctoral dissertation.

----. To appear. *A Geography of Case Semantics: The Czech Dative and the Russian Instrumental.*

Lakoff, George. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. Chicago and London: University of Chicago Press.

----, and Mark Johnson. 1980. *Metaphors We Live By*. Chicago and London: University of Chicago Press.

Langacker, Ronald W. 1982. 'Space Grammar, Analysability, and the English Passive'. *Language* 58.22-80.

-----. 1984. 'Active Zones'. *BLS* 10.172-188.

-----. l985. 'Observations and Speculations on Subjectivity', in John Haiman (ed.), *Iconicity in Syntax*, p.109-150. Amsterdam and Philadelphia: John Benjamins.

-----. 1986a. 'Abstract Motion'. *BLS* 12.455-471.

-----. 1986b. 'An Introduction to Cognitive Grammar'. *Cognitive Science* 10.1-40.

-----. 1986c. 'Settings, Participants, and Grammatical Relations'. *Proceedings of the Annual Meeting of the Pacific Linguistics Conference* 2.1-31.

-----. 1987a. *Foundations of Cognitive Grammar*, vol. 1, *Theoretical Prerequisites*. Stanford: Stanford University Press.

-----. 1987b. 'Nouns and Verbs'. *Language* 63.53-94.

-----. 1987c. 'Grammatical Ramifications of the Setting/Participant Distinction'. *BLS* 13.383-394.

-----. 1988a. 'Autonomy, Agreement, and Cognitive Grammar', in Diane Brentari *et al.* (eds.), *Agreement*

in Grammatical Theory, p.147-180. Chicago: Chicago Linguistic Society.

-----. 1988b. 'A View of Linguistic Semantics', in Rudzka-Ostyn, p.49-90.

-----. 1988c. 'A Usage-Based Model', in Rudzka-Ostyn, p.127-161.

-----. 1989. 'Absolute Construal', in F. J. Hevyaert and F. Steurs (eds.), *Worlds Behind Words: Essays in Honour of Prof. Dr. F. G. Droste on the Occasion of His Sixtieth Birthday*, p.65-75. Leuven: Leuven University Press.

-----. 1990. 'Subjectification'. *Cognitive Linguistics* 1.5-38.

-----. In press. *Concept, Image, and Symbol: The Cognitive Basis of Grammar*. Berlin: Mouton de Gruyter.

-----. To appear. *Foundations of Cognitive Grammar*, vol. 2, *Descriptive Application*. Stanford: Stanford University Press.

Lindner, Susan. 1981. *A Lexico-Semantic Analysis of English Verb-Particle Constructions with UP and OUT*. San Diego: UCSD doctoral dissertation.

-----. 1982. 'What Goes Up doesn't Necessarily Come Down: The Ins and Outs of Opposites'. *CLS* 18.305-323.

Maldonado, Ricardo. 1988. 'Energetic Reflexives in Spanish'. *BLS* 14.153-165.

Newmeyer, Frederick J. 1983. *Grammatical Theory: Its Limits and Its Possibilities.* Chicago and London: University of Chicago Press.

Poteet, Stephen. 1987. 'Paths Through Different Domains: A Cognitive Grammar Analysis of Mandarin *Dao*'. *BLS* 13.408-421.

Rice, Sally. 1987a. 'Towards a Transitive Prototype: Evidence from Some Atypical English Passives'. *BLS* 13.422-434.

-----. 1987b. *Towards a Cognitive Model of Transitivity*. San Diego: UCSD doctoral dissertation.

-----. 1988. 'Unlikely Lexical Entries'. *BLS* 14.202-212.

Rudzka-Ostyn, Brygida (ed.) 1988. *Topics in Cognitive Linguistics.* Amsterdam and Philadelphia: John Benjamins.

Smith, Michael B. 1985. 'An Analysis of German Dummy Subject Constructions'. *Proceedings of the Annual Meeting of the Pacific Linguistics Conference* 1.412-425.

-----. 1987. *The Semantics of Dative and Accusative in German: An Investigation in Cognitive Grammar*. San Diego: UCSD doctoral dissertation.

Tuggy, David. 1980. '|Ethical Dative and Possessor Omission Si, Possessor Ascension No!' *Work Papers of the Summer Institute of Linguistics, University of North Dakota* 24.97-141.

-----. 1981. *The Transitivity-Related Morphology of Tetelcingo Nahuatl: An Exploration in Space Grammar*. San Diego: UCSD doctoral dissertation.

-----. 1986. 'Noun Incorporations in Nahuatl'. *Proceedings of the Annual Meeting of the Pacific Linguistics Conference* 2.455-469.

-----. 1989. 'The Affix-Stem Distinction in Orizaba Nahuatl'. Linguistic Agency University of Duisburg.

Vandeloise, Claude. 1984. *Description of Space in French*. San Diego: UCSD doctoral dissertation.

-----. 1985a. 'Au-dela des Descriptions Geometriques et Logiques de l'Espace: Une Description Fonctionnelle'. *Lingvisticae Investigationes* 9.109-129.

-----. 1985b. 'Les Prepositions *Sur/Sous* et la Relation *Porteur/Porte*'. *Leuvense Bijdragen* 74.457-481.

-----. 1986. *L'Espace en Francais*. Paris: Editions du Seuil.

-----. 1987. 'La Preposition *a* et le Principe d'Anticipation'. *Langue Francaise* 76.77-111.