Lending

ARIEL INFORMATION:

Ariel Address: 150.135.238.50

ODYSSEY INFORMATION:

Odyssey Address: 150.135.238.6

AZN
Northern Arizona University Cline Library
Document Delivery Services
Box 6022
Building 28, Knoles Drive
Flagstaff, AZ 86011
(928) 523-6808 / (928) 523-6860 (fax)
OCLC: AZN
NAU Ariel: 134.114.228.9
document.delivery@nau.edu

TN: 191896

Borrower: AZU
Call #: PE1106.J3
Location: Book Stacks IN LIBRARY

Photocopy (Exempt)

ILL #: 37363744

Lending String: *AZN
Patron: WEIDENBACHER, HOLLIS J

Shipping Address:
ILL
UNIVERSITY ARIZONA LIBRARIES
1510 E UNIVERSITY
TUCSON AZ 85721-0055

Fax: (520) 621-4619
EMail: askill@u.library.arizona.edu

Journal Title: Readings in English transformational grammar

Article Author: Jacobs, Roderick A., comp.

Article Title: Bever, T.G., & Rosenbaum, P.; Some lexical structures and their empirical validity,

Volume:
Issue:
Month/Year: 1970
Pages: 3-19

Paged by ____ (Initials)

Reason Not Filled (check one):

☐ NOS ☐ NFAC (GIVE REASON)

☐ LACK VOLUME/ISSUE ☐ OVER 100 PAGES

☐ PAGES MISSING FROM VOLUME

Rec'd: 11/20/2007 02:05:48 PM

Some Lexical Structures and Their Empirical Validity

1

THOMAS G. BEVER
PETER S. ROSENBAUM

Discussions of a semantic theory of language can founder on the determination of the data which that theory is intended to describe. In this paper we shall avoid questions of the legitimacy of particular goals for semantic theory. Nor are we concerned primarily with the nature of the computational device used in analyzing the meanings of whole sentences. Rather we report some recent investigations of several kinds of lexical devices which are necessary formal prerequisites for the description of semantic phenomena.

The empirical reality of any descriptive device is initially supported by the facts it describes. It is further supported by the empirical validity of predictions which it justifies. In this presentation we show that the devices proposed for an adequate description of semantic lexical structure also correctly predict the potential occurrence of certain types of words and interpretations and correctly reject the potential occurrence of other types of words and interpretations.

The empirical phenomena and the formal devices which describe them involve the representation of classes of words by features, like "animateness," "plantness," and "livingness." These features are themselves organized hierarchically (e.g., anything that is "living" must be a "plant" or "animate"). Finally, certain hierarchies express relations between particular words rather

SOURCE: This work was supported by the MITRE Corporation, Bedford, Massachusetts, Harvard Society of Fellows, NDEA, A. F. 19(68)-5705 and Grant #SD-187, Rockefeller University and IBM. A preliminary version of this paper was originially written as part of a series of investigations at MITRE Corporation, Summer 1963, and delivered to the Linguistic Society of America, December 1964. We are particularily grateful to Dr. D. Walker for his support of this research and to P. Carey and G. Miller for advice on this manuscript.

than between classes of words. At each point linguistic facts force certain formal constraints on these descriptive devices.

Binary Features

Consider first the use of features in semantic description. It is not radical to propose that words are categorized in terms of general classes. Nor is it novel to consider the fact that these classes bear particular relations to each other. Consider, for example, the English words in the matrix presented in (1):

(1)	boy	tree	sheep	carcass	log	water	blood	* "og"	* "triffid"
human	+	−	−	−	−	−	−	−	−
animate	+	−	+	+	−	−	−	−	−
living	+	+	+	−	−	−	+	−	+
plant	−	+	−	−	+	−	−	+	+

On this chart there is a representation of some words in terms of particular binary dimensions. In each case the particular semantic dimension, or "semantic feature," is the marking which specifies particular aspects of the semantic patterning of the lexical item. In (2) the positive value for each feature is associated with some sample sentence frames which indicate the kind of systematic restrictions each feature imposes. Thus in (2a) any noun marked [+human], among other things, can appear in the sentence "the *noun* thought the problem over," but any noun marked [−human] cannot appear in the frame "the *noun* sensed the danger"; any noun marked [+living] can appear in the frame "the *noun* died," and any noun marked [+plant] can appear in the frame "the *noun* has damaged chlorophyll." The privileges of occurrence in these kinds of frames of the different nouns in the figure are indicated by the placement of the "+" and "−" markings on the separate features.

(2) a. [+human]: The _____ thought the problem over.
answered slowly.
was a beautiful actor.

b. [+animate]: The _____ sensed the danger.
ate the food.
ran away.

c. [+living]: The _____ died.

d. [+plant]: The _____ has damaged chlorophyll.
needed to be watered.
grew from the seed.

That such classifications play a role in natural language has rarely been questioned.[1] But the problem has been to decide which of the many possible aspects of classification are to be treated as systematically pertinent to a semantic theory. We cannot claim to have discovered all and only the semantic features of natural language. However, it is absolutely necessary to assume that there is some universal set from which particular languages draw their individual stock.

Furthermore, if the semantic analysis of natural language is to achieve explanatory adequacy, there must be a precise manner of deciding among competing semantic analyses. In this way the semantic analysis which is ultimately chosen can be said to be the result of a formal device, and not the result of the luck of a formal linguist. If the particular analysis is chosen on the basis of precise, formal criteria, then the reality of the predictions made by the device constitutes a confirmation of the general linguistic theory itself.

Lexical analysis of single words includes a specification in terms of a set of semantic features. We propose furthermore that the most highly valued semantic analysis which meets these constraints be the one which utilizes the smallest number of symbols in a particular form of semantic analysis. For a given natural language this will direct the choice of which features are drawn from the universal set as well as the assignment of predictable features in the lexicon itself.

A Hierarchy of Semantic Features

For instance, the features in the matrix in (1) are specified with some redundancy. We can see from inspection of the frames in (2a) and (2b) that if a noun can act as [+human], then it can occur as [+animate] although the reverse is not necessarily true. For example, if we can say "the man thought it over," we also can say "the man sensed the danger"; but if we can say "the ant ran away," we cannot necessarily say "the ant answered slowly."

Consequently, any word marked [+human] can be predicted as being [+animate]. This prediction, represented in rule (3a), obviates the necessity of lexical entries for "animateness" in nouns which are marked as "human."

(3) a. [+human] → [+animate]
b. [+animate] → [−plant]

Similarly, it is predictable that if a noun is [+animate], it is [−plant] in English, although the reverse is not true. Then, rule (3b) allows the features[−plant] to be left out of the lexical entry for "ant." Such redundancy rules simply reflect a class-inclusion hierarchy among the semantic features [e.g.,(3c)].

(3) c.

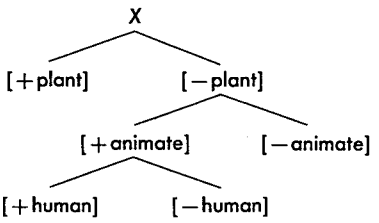

The evaluation criterion requires the simplest acceptable solution to be used. Accordingly, the solution which includes the semantic redundancy rules in (3) is clearly indicated over a solution without these rules, given that the theory allows such rules at all. [Notice that rule (3b), at least, is *not* universal, but is a rule peculiar to English, since there are languages in which at least some plants can be considered animate.]

We now can ask if the predictions made by this analysis about new lexical items are valid or not. Consider the two potential English words, "og" and "triffid," on the right-hand side of (1). "Og" is intended to be a word which signifies the dead remains of an entire plant, in the way which "carcass" indicates the dead remains of an entire animal. "Triffid," on the other hand, indicates a kind of pine tree that moves and communicates.

Neither of these words exist in everyday English, and they may both be classified as semantic lexical gaps — that is, they are combinations of semantic features which are not combined within a single lexical entry. Yet the reasons for their nonoccurrence are radically different according to the analysis in (3). Rule (3b) in principle blocks the appearance of any word like "triffid" which is both plant and animal; but the word "og" is not blocked by anything in the grammar itself. That is, "triffid" is classified by our analysis as a *systematic gap*, while the nonoccurrence of "og" is entirely *accidental.*

This distinction, made as an incidental part of our analysis, appears to be supported by the intuitions of speakers of English. Nobody would be surprised to learn that professional loggers and foresters have some word in their vocabulary which indicates the whole of a dead tree. But everybody would be surprised if they found out that foresters have a single word for trees which they think are animals. (It is that sense of surprise for English speakers which made possible *The Day of the Triffids,* a science-fiction novel in which the discovery of this concept is a primary vehicle for the plot.)

This result represents an independent corroboration of the general form of the semantic lexicon and the evaluation criteria which we proposed above. The rules in (3b) were not introduced to block systematic lexical gaps, but to achieve the highest valued semantic analysis. The fact that the systematic gaps they predict correspond to the intuitive ones is a distinct empirical corroboration of the precise formal characteristics of this form of semantic theory.

We have shown that an explanatory lexical semantic theory is possible, and have indicated some of its characteristics and empirical validations. In the next sections we briefly indicate some other semantic lexical devices, the lexical distinctions which they predict, and the validity of those predictions.

Hierarchies Among Lexical Items

Not all semantic phenomena can be handled by binary features, even if those features are related in a class-inclusion hierarchy. Consider the sentences in (4). Why is it that the sentences on the right are semantically deviant? It is not simply the case that they are counterfactual, since [4e (i)] is counterfactual, but does not involve the same kind of violation as [4e (ii)]. The first doesn't happen to be true; the second couldn't possibly be true.

(4a)	(i) The arm has an elbow.	(ii) * The elbow has an arm.
(4b)	(i) The elbow is a joint.	(ii) * The joint is an elbow.
(4c)	(i) The finger has a knuckle.	(ii) * The knuckle has a finger.
(4d)	(i) The thumb is a finger.	(ii) * The finger is a thumb.
(4e)	(i) The ant has an arm.	(ii) * The arm has an ant.

These and other similar facts justify the assumption that lexical entries themselves are arranged in two simultaneous hierarchies: One represents the inalienable inclusion of *Have*, and the other represents the class membership of *Be*. In (5) the solid vertical lines represent the *Have* hierarchy and the dotted horizontal lines represent the *Be* hierarchy. (We use the convention that the subject of an inalienable *Have* sentence must dominate the object in this tree, and that the subject of a generic *Be* must be dominated by the predicate object. Thus we can say, "the body has an arm," "the arm is a limb," but not the reverse.)

(5)

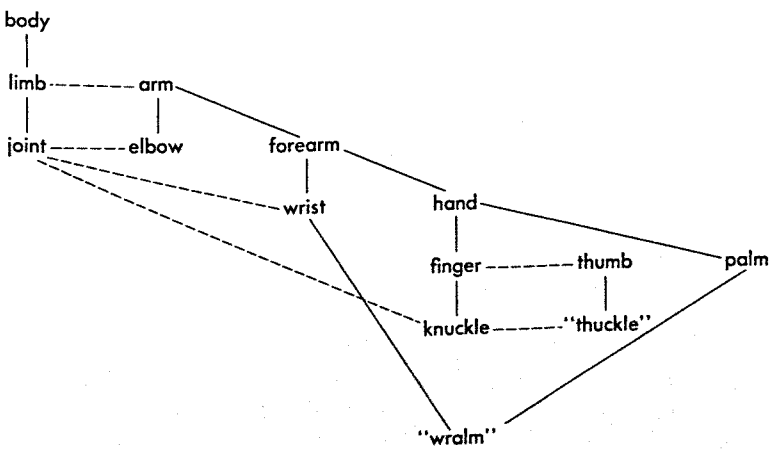

Figure (5) demonstrates why it is the case that semantic features with a limited number of states cannot handle these relations and block the kinds of incorrect sentences on the right in (4). Figure (5) allows the sentences in (6).

(6)

A body has an arm.	An arm is a limb.
An arm has a forearm.	An elbow is a joint.
A forearm has a hand.	A thumb is a finger.
A hand has a finger.	A wrist is a joint.
A finger has a knuckle.	

Any reversal of the nouns in these sentences would produce a deviant sentence, so that just to handle these cases a binary feature analysis would have to use six states.

It is clear that we must distinguish between lexically permanent hierarchies as in (5) and the productive use of "have" and "be" in actual sentences. Thus the fact that we can say "the tree is nice" and "a tree has importance" does not necessarily require a corresponding lexical hierarchy. These cases are "creative" (or alienable) uses of *Have* and *Be* rather than the "lexical" (inalienable) use exemplified in (4).

The Implications of Lexical Hierarchies for Some "Syntactic" Phenomena

Before showing the semantic empirical extensions of the lexical hierarchies, we first indicate certain obvious "syntactic" presences as part of lexical structure. (As could be expected for grammatical structures dependent on the lexicon, the distinction between "semantic" phenomena and the corresponding "syntactic" phenomena is extremely difficult and beside our point.)

Selectional Restrictions

We pointed out above that lexical items are marked as to the binary classes which they share. As demonstrated in (2) these class markings restrict the kinds of noun–verb–noun combinations that may appear within a clause. Such constraints are referred to as "selectional restrictions."

The assumption that lexical entries are organized in several hierarchies makes possible some simplifications of the representation of selectional restrictions. Consider these sentences:

(7)

The firearm shoots bullets.	* The sword shoots bullets.
The pistol shoots bullets.	* The knife shoots bullets.
The rifle shoots bullets.	* The dog shoots bullets.
The revolver shoots bullets.	
The derringer shoots bullets.	
The six-shooter shoots bullets.	

In general it is the case that if a lexical item can be selected in a construction, then everything dominated by that lexical item in the *Be* hierarchy (directly or indirectly) can also fit into that construction. For example, the hierarchy

$$\text{firearm} < \begin{matrix} \text{pistol} \\ \text{rifle} \end{matrix}$$

etc., predicts that everything generically true of "firearm" is specifically true of objects subordinate to "firearm." Notice that if a binary feature solution were sought for the above cases, there would be a separate feature (e.g., "±shoots-bullets")[2] corresponding to every level in the hierarchy which has a unique set of possible constructions.

A similar simplification is achieved by use of the inalienable *Have* hierarchy.

(8)	car	The venturi mixes gas and air.
	engine	The carburetor mixes gas and air.
	carburetor	The engine mixes gas and air.
	venturi	The car mixes gas and air.

That is, if an item in the *Have* hierarchy is a particular active construction, then the items which dominate it in the *Have* hierarchy also can fit into that construction.[3]

Feature Assimilation

There are independent reasons for postulating the existence of the *Have* and *Be* hierarchies. Consider the following sentences:

(9a)	(i)	The boy's knee aches.
	(ii)	* The statue's knee aches.
	(iii)	The statue's knee broke.
(9b)	(i)	The man's neck itches.
	(ii)	* The man's idea itches.
	(iii)	The man's idea convinced me.

The sentences in (9a) are well formed with respect to the *Have* hierarchy since both "boy" and "statue" stand in the same relation to "knee." But one observes that [9a (i)] has a conceptual status which differs significantly from sentence [9a (ii)]. The difference is a reflection of the fact that a "statue" is normally both inanimate and not living and hence cannot be subject to the set of events in which nouns that are both animate and living can participate. The descriptive problem concerns the fact that the anomalous interpretation of the phrase "knee aches" is dependent upon information which is not included in the lexical entry for the noun "knee." The information that a "knee" belongs to a noun marked either [+animate] [+living] or

[—animate] [—living] is necessarily dependent upon some mechanism which assigns these features to "knee" on the basis of the underlying phrase structure configurations. This mechanism provides a formal object which can be systematically interpreted by the semantic analysis.[4]

The mechanism which assimilates the relevant features in sentences (9a) explicitly refers only to the lexical *Have* hierarchy, rather than to any sentence with "have." This becomes clear on consideration of the fact that such assimilation does not apply in the event that the "head" noun does not stand in an inalienable relation to its modifier, as in [9b (ii)]. Underlying [9b (ii)] is the string "the man has an idea" in which the interpretation of "have" is one of "possession." As was assumed earlier, there are independent reasons for assuming that the nonlexical "have" of "possession" is distinct in certain ways from the lexical "have" of "inalienability." This hypothesis is confirmed with respect to feature assimilation. The hypothesis that feature assimilation involves the lexical *Have* hierarchy is supported further by the fact that sentences in which the *Have* hierarchy is used to its fullest extension have properties identical to those in (9a), as demonstrated in (10).

(10) a. **The toe of the foot of the leg of the boy aches.**
b. * **The toe of the foot of the leg of the statue aches.**

There is a similar additional motivation for the *Be* hierarchy as a necessary component of the grammar.

(11) a. **A gun which is a cannon fires cannonballs.**
b. * **A gun which is a six-shooter fires cannonballs.**

The descriptive problem posed by the sentences in (11) exactly parallels that posed by the sentences in (9a). The peculiarity of (11b) is reflected by the fact that only a gun which has the properties of a cannon can "fire cannonballs." The difference between the two sentences is formally characterized by an assimilation rule which assimilates the features of the modifying noun to the "head" noun: In (11) the features of "cannon" and "six-shooter" are assimilated to "gun." Such a feature assimilation provides the information requisite to the semantic interpretation of the sentences in (11) which is not systematically given by the underlying structure.[5] In particular, it is necessary to know what kind of "gun" is involved in order to determine the semantic well-formedness of the predicate.

Comparatives

There is a set of restrictions on comparative constructions which is characterized exactly in terms of the *Be* hierarchy. Consider, by way of illustration, the following sentences:

(12a) (i) **A cannon is more deadly than a pistol.**
(ii) **A pistol is more deadly than a cannon.**

(iii) * A cannon is more deadly than a gun.
(iv) * A pistol is more deadly than a gun.

(12b) (i) * A gun is more deadly than a cannon.
(ii) * A gun is more deadly than a pistol.

An examination of the compared nouns in the sentences above in terms of the *Be* hierarchy (13) reveals that the restrictions observed in these sentences can be expressed in terms of the dominance relation obtaining in this hierarchy.

(13)

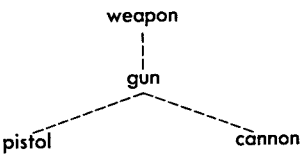

One observes that the comparative constructions in (12) are grammatical just in case a comparing noun neither dominates nor is dominated by a compared noun in the *Be* hierarchy. Thus, "pistols" and "cannons" may be compared. Furthermore, any item which "pistol" dominates (e.g., "derringer") can be compared with either "cannon" or any other item dominated by "gun" which neither dominates nor is dominated by "pistol" (e.g., "rifle") as in the following examples:

(14) a. A cannon is more deadly than a derringer.
b. A rifle is more deadly than both a derringer and a cannon.

Although the derivation of comparative constructions in general has not been fully resolved, it is nonetheless clear that certain restrictions on these constructions require reference to a lexical *Be* hierarchy.

Hierarchies Are Nonconvergent

The examples we have discussed do not involve convergence within one hierarchy. That is, there are no instances of hierarchies like (15)

(15)

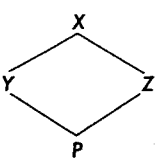

in which the hierarchy branches at *X* and converges at *P*. There are several reasons why convergent hierarchies are not tolerable. First, there are many instances in which a single lexical item is dominated by different words, e.g., (16).

(16) humans — hands (excluding intermediates) clocks — hands

If convergences were an acceptable part of the formalism, it would be possible to simplify (16) to (17).

(17) humans, clocks → hands

But this would transfer to "hands" all the selectional features of both "humans" and "clocks," many of which are mutually incompatible (notably [± animate]). Furthermore, if the senses of "hands" are not distinguished, we could not account for the oddity of sentences like: "The clock's hands look funny, but mine don't."

Maintaining distinct senses of lexical items in distinct lexical hierarchies is motivated independently of the consequences for convergent hierarchies. There are also motivations against convergent hierarchies within the same "lexical tree." Consider the subsection of Figure (5) in (18).

(18) limb — arm; joint — elbow; joint — knee

"Arm" could dominate "joint" directly (as well as via "limb") since there is a sentence "the arm has a joint." However, this would transfer to "joint" all the properties pertaining to "arm." In turn this would transfer all the features of "arm" to "knee," which is clearly incorrect.

The requirement that there are no convergences in the lexical hierarchy has the result that there are many distinct matrices in the lexicon which utilize the *Have* and *Be* hierarchies.[6] There are many problems about these matrices which require further investigation. We have merely outlined the essential motivations for the hierarchies as lexical structures.

Systematic and Accidental Lexical Gaps Explained by Lexical Hierarchies

The evidence above attests to the general need for lexical hierarchies as one of the structures of grammar. The hierarchies provide mechanisms for the blocking of certain obvious anomalies (e.g., [4a (ii)]) and for the simplification of general selectional processes. The hierarchies also make predictions about the kinds of lexical items which can and cannot occur. That is, they provide a further basis for distinguishing lexical gaps which are systematically ex-

plained by the lexical structure of the language from those gaps which have no such explanation and are therefore "accidental."

For example, the combination of the hierarchies as in (5) makes several claims about English lexical structure. The placement of "finger," "thumb," and "knuckle" in the lower right of (5) strongly implies that the single lexical item indicating a "thumb knuckle" would be a regular extension of the present lexical system, since it would fit in the regular slot (labeled "thuckle"). An instance of an irregular extension would be the word which is a unique part of both the wrist and the palm [labeled "wralm" in Figure (5)]. The irregularity of "wralm" is explained by the fact that it would require a convergence of the hierarchy, which is not allowed in our formulation, for independent reasons. Thus, "wralm" would require two separate hierarchies and would involve added complexity.

Surely this result is an intuitive one — the concept "thumb knuckle" seems quite natural, but the concept of "wrist palm" does not. It should be reemphasized here that it is irrelevant to us whether or not these linguistic properties are reflections of something in the so-called "real world." Our present quest does not include the source for linguistic structure but seeks to define the nature of the structure itself.

The hierarchy may also predict instances in which certain uses of words are dropped from a language. Consider the treatment of the word "sidearm" in the *Be* hierarchy introduced above:

A $\begin{Bmatrix}\text{pistol}\\ \text{sword}\end{Bmatrix}$ is a sidearm.

A sidearm is a weapon.

These sentences indicate that "sidearm" would have to be represented with a convergence on "pistol" as in (19), or with two lexical instances of "pistol" as in (20).

(19)

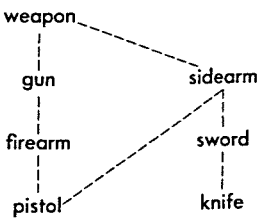

(20)

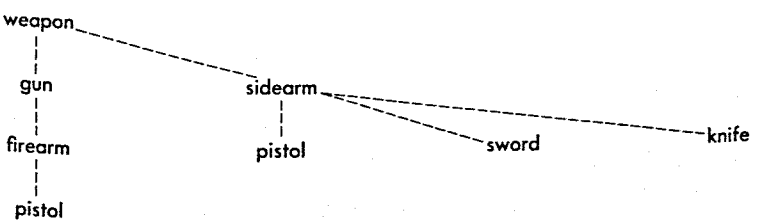

If it were not for the fact that swords (and knives, technically) are sidearms, it would be possible to formulate the hierarchy without any double dominance relations:

(21)


In fact, although we must allow for the possibility of the facts represented in (19), we have set up the formalism such that that type of structure involves a large amount of complexity. Example (20) is represented as two separate hierarchies to avoid the double dominance of "pistol." This offers an explanation of the tendency for the term *sidearm* (meaning "sword" and "knife") to drop out of modern English, since it reduces the complexity of (20) to (21).

As above, we cannot rule out any particular aggregate of features (or possible constructions) on the basis of lexical structure rules, since anything *can* occur as a lexical item. But the theory which we set up to represent lexical structure provides some formal basis for distinctions between lexical items which do occur. For instance, "shotgun" is in an anomalous position in the hierarchy because of these sentences:

(22) a. **A shotgun is a firearm.**
b. **A firearm shoots bullets.**
c. *** A shotgun shoots bullets.**

This requires us to mark "shotgun" as an exception to the hierarchy (or to the rule which predicts features on the basis of the hierarchy).

In certain cases the restriction on convergence in the hierarchy is supported directly by intuitions. For example, consider the word "neck:"

(23) a. **John's head has a thick neck.**
John's head was frozen, including his neck.

b. **John's torso (body) has a thick neck.**
John's torso (body) was frozen, including his neck.

For some people, sentences (23a) are correct and (23b) are incorrect; for others the reverse is true. From many, both (23a) and (23b) are acceptable at different times, *but not simultaneously:* that is, there is no single lexical hierarchy which includes a structure like (24),

(24)

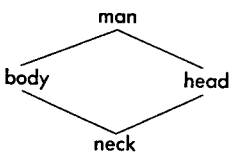

but rather the hierarchy is like (25).

(25)

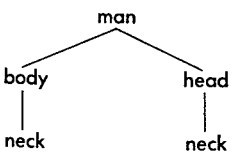

It is tempting to "explain" these facts as an "obvious" reflection of the fact that the neck resides between the body and the head. But there is nothing incompatible with the physical facts in an analysis which would delegate the neck as part of the "head" and simultaneously part of the "body." Thus, it is not "reality" which disallows convergences of lexical hierarchies, but the properties of the lexical system itself.

Metaphor and the Form of Lexical Entry

The final case we shall consider involves what is sometimes called "historical metaphor." For instance, in modern English there is a regular rule [see (26)] which extends surface quality adjectives which are drawn from a restricted set of abstract qualities. Thus, in (27b) the adjective "colorful" is used with the noun "ball" or "idea." Certain adjectives drawn from a restricted set are not affected by rule (26). For example, shapes and colors

(26)

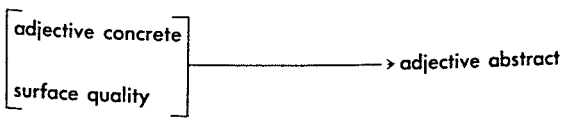

are not regularly applied to ideas, as can be seen in (27d). Whenever shapes are used with abstract nouns, as in (27e), they are defined by our analysis as unique expressions and not as the result of the general rule in (26).

(27) a. **The ball is colorful.**
b. **The idea is colorful.**

c. **The ball is ovoid.**
d. *** The idea is ovoid.**

e. **The idea is square.**
f. **The idea is variegated.**

This too seems intuitively correct. In "the idea is square" we know that an interpretation is possible but unique. In this case we happen to know the unique abstract interpretation of "square" but not in "the idea is ovoid." Although we may never have heard the sentence "the idea is variegated" before, we can immediately recognize it as a regular lexical extension of the "literal" meaning of "variegated."

The representation of lexical items in hierarchies facilitates the automatic interpretation of the metaphorical extensions of certain words from their original lexical structure. Consider the analysis of the word "colorful," using the *Be* and *Have* hierarchies simultaneously.

(28)

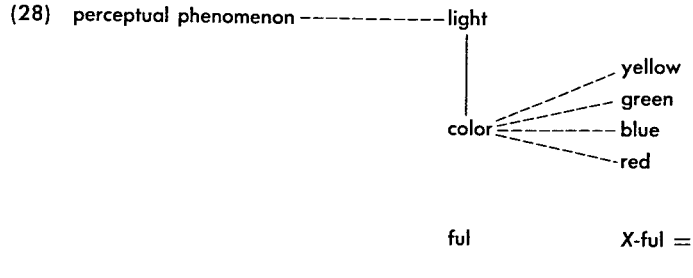

Suppose we use the following metaphorical principle: A metaphorical extension of X includes the hierarchical relations of the literal interpretation without the specific labels: i.e., in metaphorical extensions (28) becomes (29).

(29)

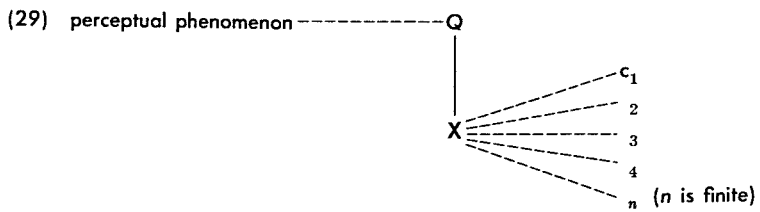

In this way the word "colorful" as applied to "idea" is interpreted as "having simultaneously a large set of potentially distinguishing characteristics." Similarly for "shape" in "the idea took shape."

(30)

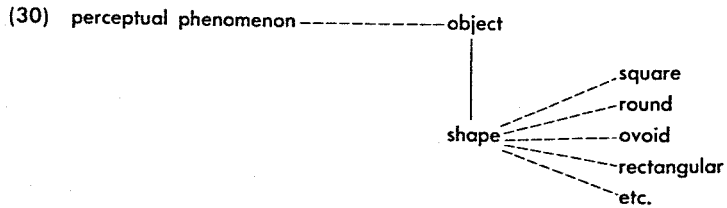

In this usage, "shape" is interpreted as "a particular character out of many (but finite) characters." (Note that although there is an infinite variety of

shapes possible, there is a finite number listed in the lexicon; what "the idea has shape" means is that it has one clear, recognizable characteristic, as opposed to an unnameable one.) Consider now the interpretation of "the idea is square." It is "the idea has a specific characteristic out of a finite list of them." But there is no special indication as to which characteristic it is, or how it fits into the abstract realm of ideas. Thus it takes special knowledge to interpret the metaphorical extensions of lexical items at the most subordinate part of the hierarchy, if it takes special knowledge in the original literal interpretation.

For these kinds of cases this principle provides a more general rationale for metaphorical interpretation than does rule (26). It distinguishes cases of metaphor which require no special knowledge ("shape, color") from those which do ("square, green"). It also offers some insight into the nature of the metaphorical extension itself.

There are many metaphors for which this sort of lexical account will not do. Thus, "the secretary of defense was an eagle gripping the arrows of war, but refusing to loose them" is not interpretable in terms of any lexical hierarchy. What we have shown is that the interpretation of "lexical metaphors" can proceed in terms of the original form of the lexical entry, and that this can distinguish between "regular" metaphorical extensions and isolated cases.

Conclusion

The force of the empirical validity of lexical distinctions and the corresponding descriptive mechanisms depends on the extent to which the analyses are arrived at according to a formal evaluation criterion. The exact form of the distinctive features in (1), the feature hierarchy in (3c), the lexical hierarchy in (5), and the metaphor rules will ultimately be determined by the form of universal semantic structures and the complexity of a particular description.

The descriptive devices we have presented do not exhaust what is necessary for semantic theory. What we have shown is that an explanatory semantic theory is possible, and we have indicated some of its features. Primarily, we have tried to show what kinds of empirical considerations must be included in the evaluation of different candidates for an explanatory semantic theory.

NOTES

1. Modern discussions of such features and their integration within grammatical theory can be found in Katz and Fodor (1963); Chomsky (1965); Miller (1967). In this article we will assume that the reader has a basic familiarity with the role of semantic analysis in current transformational linguistic theory.

2. We are not claiming that there are *no* features which are pertinent to particular restricted sets of lexical items — in fact, if one argued that the feature [± shoots

bullets] should be used, then our argument simply is that the feature [+shoots bullets] is predictable for everything below its first occurrence in the *Be* hierarchy. In other words, if other aspects of the formal treatment of the above problem require features for uniformity of notation, this can be accommodated easily. Nevertheless, it remains the case that the hierarchy can be utilized to reduce intuitively the duplication of lexical information.

3. Note that there is something odd about "the car mixes gas and air" although it is technically correct. There are other cases like this: "The electric lamp has tungsten," "The body has fingernails." There are several potential explanations for the oddness of these sentences: (a) Certain words (e.g., "carburetor") are designated as referring to "the whole" of an object and feature assimilation cannot pass through them. (b) We must distinguish between various senses of *Have:* "have *in* it." "have as *part of* it," "have *adjacent to* it." Then a car might be said to have a "carburetor" *in* it but not as *part of* it, while the "venturi" is *part of* the carburetor. It would not be the case that a car and a carburetor "have" a venturi in the same sense of *Have.* (c) There might be a principle of linguistic performance: The more nodes an assimilation of features passes through, the lower the acceptability of the sentence.

4. It is conceivable that the assimilation rules are to be incorporated into the interpretative semantic cycle itself rather than to provide a derived semantic structure upon which the interpretative component operates. Factors bearing on this decision will not be explored in this paper. At the time when we first proposed the problem raised by sentences like those in (9), several of our friends and teachers argued that a syntactic solution would be forthcoming. This has been incorporated as part of the analysis of adverbs proposed by G. Lakoff. Briefly, [9a(i) and (ii)] would be analyzed as "the boy aches in his leg" and "the statue aches in his leg." The anomaly of the second would be explained as a function of the fact that "the statue aches" is anomalous. Sentence [9b (ii)] would be anomalous because there is no sentence "the man itched in his idea," while there is a sentence corresponding to [9b (i)] "the man itched in his neck." Whether or not one is convinced by this analysis (which we are not), the fact remains that those nouns which can act as agents with verbs, like "ache," "itch," "hurt," etc., are just those which are represented in the lexical *Have* hierarchy as inalienably part of the "head" noun. For example, there is a sentence "I hurt" and "I hurt in my arm," but no "I hurt in my carburetor." Similarly, words like "idea" or "thought" are not entered lexically. So either form of syntactic analysis of these examples presupposes a lexical *Have* hierarchy. (Note that this does not preclude a "creative" grammatical component which would allow for new *ad hoc Have* and *Be* relations to be produced as needed. As we continually emphasize, the ontogenesis of lexical structures is not of concern to us in this paper.)

5. In fact, it might appear that an incorrect result follows from this restriction just in case a particular property or object is shared by many different words. For example, consider the word "electron." Convergent hierarchies would allow this word to appear only once in the lexicon:

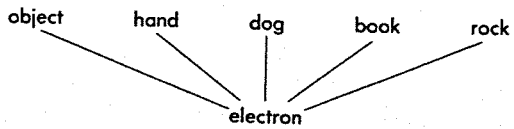

The restriction against convergent hierarchies would appear to force such a word to have multiple representation:

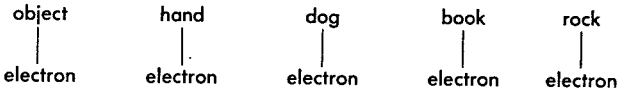

However, "electron" does not appear many times, once under each object; rather only once under the word "object" itself. Any other word that *is* an object automatically acquires the property of *having* electrons, by the assimilation rules discussed above.

6. These sentences seem to suggest that sentences like "a gun is a cannon" must be generated in the grammar in order to account for the relative clause formation. In other words, the original motivation for the lexical *Be* hierarchy would seem to disappear. This argument is untenable, however, since it is necessary to postulate that a relativized noun is definite, referring to a *specific* noun: e.g., "A gun fires cannonballs if it (that gun) is a cannon."