THE SEMANTICS OF SENTENCE SUBJECTS*

HERBERT H. CLARK** and JEFFREY S. BEGUN
Carnegie-Mellon University

It is proposed that subjects of transitive verbs in English are perceived as belonging to a semantic hierarchy, with human subjects most acceptable and non-human subjects less so. As a test, six semantically defined classes of subjects were systematically interchanged in simple, well-formed sentences: in Experiment I, 25 students rated the resultant sentences for acceptability; in Experiment II, 25 more students altered one word in the same sentences to make more sense. The results affirmed the following hierarchy: human nouns, animal nouns, concrete-count nouns, concrete-mass nouns, abstract-count nouns, and abstract-mass nouns. In Experiment I, a subject higher in the hierarchy was found to replace one lower down more sensibly than the reverse; in Experiment II, the alterations were consistent with these judgments. The results suggest that, linguistically, the features [+*Human*], [+*Animate*], and [+*Concrete*] are canonical or unmarked in the semantic representation of subjects. The consequences of this are discussed for the processes of interpreting and composing sentences.

The logical subject of the transitive verb in English can be of many semantic types. It can be human (*John watched an aardvark*), non-human but animate (*The aardvark was eating his lunch*), concrete but inanimate (*The lunch pleased the aardvark very much*), abstract (*His ugliness frightened John away*), and so on. Yet one of the most common observations about subjects in English (see, e.g., Clark, 1965; Svartvik, 1966) is that they tend to be human or at least animate. This observation has been so compelling that many linguists have tried to characterize the subject-predicate relation semantically as actor and action: in *John watched the aardvark,* John is the actor who carries out the action of watching the aardvark. But this characterization, as Chomsky (1957) rightly observed, is not general, for *lunch* in *The lunch pleased the aardvark very much* cannot be considered an "actor" or "doer" in any sense of these words. How, then, are we to reconcile these two observations—that subjects tend to be actor-like, yet can be otherwise? Or, more generally, what is the English speaker's conception of the subject of the transitive verb? The present experiments were designed to give at least partial answers to these two questions.

The study of subjects is complicated, however, by the fact that it is inextricably tied to the study of verbs. It is convenient, then, to discuss subjects and verbs together by introducing the linguistic notion of selectional restrictions. These, according to Chomsky (1965), are the rules which serve to delimit a verb's use with subjects and objects. *Read,* for example, obeys a rule restricting its use to animate subjects: to say *The map read the book* is nonsensical, but to say *The man read the book* is not. The prevalence of animate subjects in English can therefore be described in terms of

* *This research was supported in part by Public Health Service research grant MH-07722 from the National Institute of Metal Health. We thank Margaret Poll for her assistance.*

** *Now at Stanford University.*

selectional restrictions on verbs. Imagine the situation if all subjects had to be animate in English; under these conditions, all transitive verbs would have to have selectional restrictions making them unacceptable to all but animate subjects. But in fact, there is only a *preference* for animate subjects in English. This might have one of the three following consequences on the system of verbs: (1) Verbs could belong to mutually exclusive categories. Because of selectional restrictions, some verbs would be acceptable only with human subjects, some with inanimate concrete subjects, some with abstract subjects; and so on. But, overall, there would be many more verbs which are acceptable with animate than with non-animate subjects. (2) Verbs could belong to hierarchial categories. All verbs would be acceptable with human subjects; a sub-set of these would be acceptable with inanimate concrete subjects; a sub-set of the latter would be acceptable with abstract subjects; and so on. This scheme, of course, would also arrange subjects hierarchically, with human subjects on top and the other types successively lower down. Or (3) there could be a mixture of (1) and (2). We might call case (1) the Exclusive case, (2) the Hierarchical case, and (3) the Mixed case.

There is some evidence available that supports the Hierarchical over the Exclusive case, although the true situation is almost certain to be a mixture of the two. Notice, for example, that most sentences with abstract subjects (like *Sincerity frightened John*) are equally acceptable with a substituted animate subjects (*Mary frightened John*), although many sentences with animate subjects (*Mary ate lunch*) are not acceptable with a substituted abstract subject (*Sincerity ate lunch*). This suggests the Hierarchical case in which human and abstract subjects are arranged, not as two mutually exclusive classes, but as two hierarchically arranged classes with the human class on top. Also, in a previous experiment (Clark and Begun, 1968), it was found that subjects and objects were asymmetrical in just the same way: a sentence was judged more acceptable when its subject had been replaced by the subject, rather than by the object, of another sentence. This indicated that, as the subject of a sentence, other subjects, which were mainly animate, were arranged hierarchically above other objects, which were mainly inanimate.

Two experiments were therefore designed to explore the possibility that sentence subjects are arranged hierarchically by their semantic properties. The approach was distributional, as in Clark and Begun (1968). Nouns with various specified semantic features—like being human or animate or concrete—were systematically substituted for each other in well-formed sentences. In the first experiment, students rated the sensibleness of the resulting sentences; in the second, other students altered the resulting sentences so that they made more sense. Sentences containing a subject higher in the hierarchy that was substituted for one lower in the hierarchy were expected to be more acceptable than sentences with the reverse condition; also, subjects in the former sentences were expected to be altered less often than subjects in the latter.

Experiment I

Method

A large number of sentences, to be judged for sensibleness, were constructed in three steps: first, a sample of appropriate nouns was selected; second, students were asked to compose sentences using these nouns as subjects; and third, the parts of the composed sentences were systematically recombined to generate anomalous sentences to be judged for their sensibleness.

As a first step, five types of noun were chosen in accordance with Chomsky's (1965) syntactic features for common nouns. (Proper nouns were ignored in this study.) Forty instances of each noun type were sampled in a quasi-random fashion from the Thorndike-Lorge (1944) word list, with the restriction that each word had a frequency count of over 50 per million on the general count. Nouns were classified into five types by the following criteria: (1) *Human* nouns had human referents and could be pluralized; (2) *Animal* nouns had non-human animate referents and could be pluralized; (3) *Non-animate-count* nouns had non-animate referents, and could be pluralized, but could not be modified by *much;* (4) *Concrete-mass* nouns had concrete referents and could be modified by *much;* and (5) *Abstract-mass* nouns had non-concrete referents and could be modified by *much.* There were not enough instances of Animal nouns in the Thorndike-Lorge lists using the frequency criterion of 50 per million, so for this type the criterion was lowered to 25 per million.

Table 1

Syntactic features of five types of noun.

	SYNTACTIC FEATURES			
NOUN TYPE	Human	Animate	Count	Concrete
(1) Human nouns	+	+	+	
(2) Animal nouns	−	+	+	
(3) Non-animate-count nouns			+	
(4) Concrete-mass nouns			−	+
(5) Abstract-mass nouns			−	−

Table 1 lists the five types of noun and their syntactic features from Chomsky (1965). Animal nouns, for instance, have features [−*Human*], [+*Animate*], and [+*Count*], while they are unspecified on the feature *Concrete*. The feature *Concrete* has been used in place of *Abstract,* which Chomsky uses, for purposes which will be clearer later on. Our [+*Concrete*] is equivalent to his [−*Abstract*]. Examples of the five types of noun, respectively, are: *man, teacher; dog, spider; book, fact; grain, snow;* and *harm, growth.*

The second step was to have students compose sentences using each noun as the subject of a transitive verb. For this purpose each of the 200 nouns was placed in the subject slot of the sentence frame, The ——— ———ed the ———. The definite

article *the* preceding Abstract-mass nouns, however, was omitted everywhere in the study. The sentence frames, complete with subjects, were each printed on blank IBM cards. Fifteen students were then given 40 sentence frames each, eight frames of each type of noun, and were asked to complete them as they chose with good English sentences. They were further told that the *-ed* at the end of the verb slot meant only that they should use a past tense verb. One sub-group of five students, therefore, provided sentences for all 200 noun instances, and the other two sub-groups provided an additional two sentences for each noun. An outside judge was asked to choose the " best or most sensible " sentence from the three which had been composed for each of the 200 nouns; sentences which could not be passivized had first been eliminated, The 200 sentences chosen constituted the well-formed English sentences for use in the next step.

The third step was to construct semantically anomalous sentences for other judges to rate for sensibleness. An iterative scheme was used to generate 5000 " derived sentences " from the 200 composed sentences. The derived sentences were classifiable into 25 conditions (200 derived sentences per condition) which were defined by the nature of the subjects and predicates of the sentences. To indicate the conditions, we first define a Human subject as a Human noun used as the subject of a sentence; a Human predicate, on the other hand, is defined as a predicate of a sentence composed with a Human noun as subject. The 25 conditions made up a 5×5 design in which the five types of subject were each used with the five types of predicate. One condition, for example, contained sentences all of which had Animal subjects and Concrete-mass predicates. There was the further restriction, however, that a sentence with, say, an Animal subject and an Animal predicate was never constructed from the subject and predicate from the same original sentence.

To generate all 5000 derived sentences, the 200 composed sentences were divided into eight groups of 25. Each group contained five randomly chosen sentences composed with each type of subject. The 25 subjects of one group of sentences were then assigned to the 25 predicates of another group of sentences in a 25×25 Latin square design. This Latin square consisted of the 25 subjects, the 25 predicates, and 25 judges. Each judge (other students to be described below) received 25 sentences from this square, one from each of the 25 subject-predicate conditions. A property of this design is that each subject was paired with each predicate only once; this produced 652 uniquely derived sentences. The 25 sentences any one judge received, moreover, contained all 25 subjects and 25 predicates, each used once. The same procedure was followed for the other seven groups of composed sentences, so that each judge was asked to rate eight times 25, or 200, derived sentences, eight for each subject-predicate condition. This scheme generated, in all, 5000 uniquely derived sentences.

The procedure for judging the anomalous sentences was as follows: The 25 judges, all Carnegie-Mellon University undergraduates from Introductory Psychology classes, were each given 200 derived sentences, which had been printed, one each, on IBM cards. The judges were asked to rate each sentence for sensibleness on a 1-to-7 rating scale, on which "1" meant " very nonsensical " and "7" meant " very sensible " They

were required to write that rating directly on the data card and were urged to work quickly but accurately. The 200 sentences for each judge had been randomly ordered by shuffling the data cards.

Results

The mean ratings for the 25 subject-predicate conditions are shown in Table 2. An analysis of variance of the 5000 ratings shows: (a) the mean ratings for each subject type differed significantly, $F(4,576) = 153.06$, $p < 0.001$; (b) the mean ratings for each type of predicate differed significantly, $F(4,576) = 49.46$, $p < 0.001$; and (c) the interaction of subject and predicate was significant, $F(16,576) = 23.33$, $p < 0.001$.

TABLE 2

Mean ratings for 25 Subject-Predicate Conditions.

	TYPE OF PREDICATE					
TYPE OF SUBJECT	(1)	(2)	(3)	(4)	(5)	Means
(1) Human	6.08	4.98	4.92	4.94	5.52	5.29
(2) Animal	3.24	4.75	3.76	3.66	3.78	3.84
(3) Non-animate-count	2.48	2.42	3.82	4.03	4.11	3.37
(4) Concrete-mass	2.14	2.36	3.94	4.28	4.14	3.37
(5) Abstract-mass	2.94	2.60	3.43	3.28	4.93	3.43
Means	3.38	3.42	3.98	4.04	4.50	3.86

First, some types of subject were generally more acceptable than others. Sentences with Human subjects were regarded by the judges as very sensible, and this was true regardless of the predicate. Sentences with Animal subjects were next most sensible, and sentences with other types of subject were all less sensible but about equal. According to Duncan's multiple range test, the differences between Human and Animal subjects and between Animal and the other three subject types were significant at $p < 0.01$; the latter three subject types did not differ significantly.

The reverse was true of the predicates: Human predicates were least sensible, and Abstract-mass predicates most sensible, with the other three types of predicate falling correspondingly in between the extremes. According to Duncan's multiple range test, Human and Animal predicates did not differ from each other significantly, nor did Non-animate-count and Concrete-mass predicates; but all other differences were significant at $p < 0.01$. Thus it is the predicates composed for Abstract-mass subjects that allow the greatest variety of noun types to be substituted in the subject position. The predicates of Human and Animal subjects are by far the most restrictive.

The most important property of the matrix in Table 2, however, is that it is perfectly asymmetrical. Each cell above the main diagonal is larger than its corresponding cell below the diagonal. Consider, for example, the four cells of the Human and Animal subjects and predicates in the upper left-hand corner of the matrix. The two

main diagonal cells show how substitutable two Human nouns are for each other and how substitutable two Animal nouns are for each other. The two off-diagonal cells show how substitutable a Human noun is for an Animal noun and vice versa. But these two off-diagonal cells are asymmetrical: Human nouns fit more sensibly with Animal predicates than Animal nouns fit with Human predicates. This property will be called dominance. Human nouns will be said to dominate Animal nouns, since Human nouns make better subjects of Animal predicates than Animal subjects make of Human predicates.

From the asymmetry of the matrix, it is clear that the five types of subject form a strict dominance hierarchy. Human subjects dominate all others; Animal subjects dominate all other but Human ones; Non-animate-count subjects dominate Concrete-mass and Abstract-mass subjects; Concrete-mass subjects dominate only Abstract-mass subjects; and finally, Abstract-mass subjects dominate nothing. This ordering places Human subjects on top of the dominance hierarchy and other subjects successively lower depending on how distant from " being human " they are. This property will be considered in detail in the Discussion. According to Duncan's multiple range test, all ten asymmetries are significantly large (at $p < 0.05$), except for that between Non-animate-count and Concrete-mass subjects and predicates.

Non-animate-count subjects, however, seemed to form a fairly heterogeneous class of subjects, since the condition with Non-animate-count subjects *and* predicates had a mean rating, 3.82, much lower than the other four comparable ratings (on the main diagonal). One Non-animate-count subject was not well suited as a substitute for another such subject. One reason for the heterogeneity of such nouns is that they seem to be of two different types: concrete and abstract. Words like *tree* and *map* are concrete, but words like *fact* and *week* are abstract. This suggested a further analysis in which the two sub-classes of Non-animate-count nouns were treated separately. This division, it should be noted, is not part of the syntactic feature system in Chomsky (1965) and implies a slightly different feature system from Chomsky's.

To accomplish this, we divided the class into concrete and abstract nouns. Concrete-count nouns were defined as those with referents which could be pointed at; a map and a tree can be pointed at. Abstract-count nouns were considered those lacking this property; one cannot point at a fact or a week. There were 27 Concrete-count nouns and 13 Abstract-count nouns by this division. The predicates associated with the Concrete-count and Abstract-count nouns, of course, were also divided out. We then constructed a 6×6 matrix of mean ratings, just like the 5×5 matrix shown in Table 2. First, however, the rating of each of the 5000 rated sentences was corrected so that the mean rating for each rater was equal to the grand mean. This correction was necessary since in the 6×6 matrix there is a slight confounding of raters with the 36 conditions; the correction effectively eliminates that confounding.

The 6×6 matrix is shown in Table 3. This matrix, once again, shows a strict dominance among the noun types. In this hierarchy the ordering is: Human nouns, Animal nouns, Concrete-count nouns, Concrete-mass nouns, Abstract-count nouns,

TABLE 3

Mean ratings for 36 Subject-Predicate Conditions.

	TYPE OF PREDICATE						
TYPE OF SUBJECT	(1)	(2)	(3a)	(4)	(3b)	(5)	Means
(1) Human	6.08	4.98	4.58	4.94	5.64	5.52	5.29
(2) Animal	3.24	4.75	3.26	3.66	4.80	3.78	3.84
(3a) Concrete-count	2.65	2.49	3.58	4.09	4.31	3.72	3.28
(4) Concrete-mass	2.14	2.36	3.18	4.28	4.66	4.14	3.37
(3b) Abstract-count	2.51	2.28	3.45	3.91	5.83	4.92	3.55
(5) Abstract-mass	2.94	2.60	3.34	3.28	3.94	4.93	3.43
Means	3.38	3.42	3.60	4.04	4.76	4.50	3.86

and Abstract-mass nouns. There are no exceptions in this hierarchy; each type of noun dominates all those below it.

One final analysis shows more clearly the semantic relations among the six types of noun. Since some types appear more generally useful as subjects than others, and since some types of predicate make more sense no matter what the subject, one would like to discover if there is any structure once these two effects have been taken out of the data. To do this, we first calculated the interaction effects in Table 3 by subtracting out the differences in the row and column means. These effects seemed to contain only minor asymmetries about the diagonal. The bottom half of the matrix was therefore combined with the top half, forming a half-matrix of similarity measures. The more interchangeable, say, Human and Animal subjects and predicates are, the higher will be the combined interaction score of the two categories. This is what is meant by "similarity." Finally, a hierarchical clustering scheme—Johnson's (1967) diameter method—was used to picture the relations among the subject types. This scheme, by placing two types or clusters of types together when they have a high similarity measure, derives a hierarchy of clusters representing the matrix of similarity measures.

Significantly, the clustering that emerges from this analysis makes divisions which conform to distinctions made in the syntactic features. At the most general level of clustering, there is a distinction of animate from non-animate nouns; Human and Animal nouns fall into one cluster, and the rest into another. The next level of clustering divides the non-animate nouns into two clusters. In one cluster are the concrete nouns (Concrete-count and Concrete-mass), and in the other are the abstract nouns. Human and Animal nouns are also distinguished at about this level. It is only within the concrete and abstract clusters that count and mass nouns are finally differentiated. The property of count versus mass appears to be quite unimportant for distinguishing classes of sentence subjects.

Experiment II

Method

Twenty-five more students were given decks of derived sentences exactly duplicating the decks given to the 25 Experiment I judges. These students were instructed to read through each sentence and to change one word in the sentence—except *the*—so that the sentence made good sense. In the instructions it was stressed that even in sentences that already made sense they should alter one word so that it made still more sense. If in changing a word they also wanted to add or take out a *the*, they were told, they could do so. This allowed the Abstract-mass and other noun types to be interchangeable. These students, also Carnegie-Mellon University undergraduates from an Introductory Psychology course, were run in groups of five to ten. They completed the task in thirty to forty minutes.

Results

The subjects of less acceptable sentences were expected to be altered more often than those of more acceptable ones. As a test, the percentages of sentence subjects altered were calculated for each of the 25 subject-predicate conditions. The expectations were clearly upheld, for the rank-order correlation between those percentages and the mean ratings of the 25 subject-predicate conditions in Table 2 is -0.82, significant at $p < 0.001$.

Again a second analysis seemed to be in order. As in Experiment I, the 5×5 matrix was expanded into a 6×6 matrix, by dividing the Non-animate-count subjects and predicates into Concrete-count and Abstract-count types. In this case, nothing could be done about the confounding of students with the 36 conditions. The percentages of altered subjects in Table 4 also showed a high negative correlation with the mean ratings in Table 3, a rank-order correlation of -0.78, $p < 0.001$.

Table 4

Percentage of times subjects were altered for 36 Subject-Predicate Conditions.

	TYPES OF PREDICATES						
TYPES OF SUBJECTS	(1)	(2)	(3a)	(4)	(3b)	(5)	Means
(1) Human	62	60	47	59	55	48	56
(2) Animal	77	60	71	64	51	56	64
(3a) Concrete-count	89	84	75	79	72	64	78
(4) Concrete-mass	88	82	63	70	78	68	75
(3b) Abstract-count	86	88	74	74	58	69	77
(5) Abstract-mass	82	78	80	76	68	67	76
Means	79	73	62	69	64	61	70

The percentage data in this experiment reflect the dominance hierarchy found in Experiment I, with a few minor exceptions. In this experiment, dominance can be defined as a difference in percentages in the appropriate cells, rather than a difference

in ratings. When Animal subjects are altered more often in the context of Human predicates than Human subjects are with Animal predicates, then Human nouns are said to dominate Animal nouns. The percentages in Table 4 reflect the following dominance hierarchy: Human nouns, Animal nouns, Concrete-mass nouns, Concrete-count nouns, Abstract-mass nouns, and Abstract-count nouns. This hierarchy is perfect except for one four-percentage-point difference. The dominance order agrees with that found in Experiment I except for the reversal of count and mass nouns within the concrete nouns and within the abstract nouns.

The particular substitutions students made in altering the sentences are consonant with the five-member hierarchical ordering found in Experiment I. Each substitution was classified as one of the five types of noun originally defined in this study. Generally, there was a tendency to move up the hierarchy—to replace the given subject with nouns more closely related to human nouns. After all alterations had been taken into account, 38% of all 5000 sentences had Human subjects, 20% Animal subjects, 20% Non-animate-count subjects, 12% Concrete-mass subjects, and 10% Abstract-mass subjects. The percentages for each noun type before alterations, of course, were equal at 20%.

As expected, the substitutions were more influenced by the predicates they had to agree with than by the subjects they were replacing. For example, the subjects of sentences with Human predicates were very often changed to agree in kind with the predicates; i.e., they were changed to Human nouns. After all alterations, 81% of the sentences with Human predicates contained Human subjects. The analogous percentages for Animal, Non-animate-count, Concrete-mass, and Abstract-mass predicates and subjects were, respectively, 65%, 42%, 33%, and 20%. When the changed subject was not in agreement with the predicate—e.g., a Human subject was substituted in a sentence with an Animal predicate—the substitution was more often a noun type higher, rather than lower, in the hierarchy, by a ratio of 4.05 to 1. The hierarchy found in Experiment I is therefore very potent in predicting how a sentence will be changed to make more sense.

Discussion

The six types of subject studied in these experiments formed a hierarchy. Human nouns made the most acceptable subjects, and others made less acceptable subjects depending on how "non-human" they were. The hierarchical ordering from most to least acceptable was as follows: Human nouns, Animal nouns, Concrete-count nouns, Concrete-mass nouns, Abstract-count nouns, and Abstract-mass nouns. Experiment I showed that Human subjects could substitute for each subject type lower in the hierarchy more sensibly than the reverse; in the terminology of this study, Human subjects dominated all other subjects. More generally, each subject type dominated all other types lower in the hierarchy. Experiment II showed that Human subjects,

when substituted for other subjects lower in the hierarchy, were altered less often than when the reverse was true. This relation was also consistent down the hierarchy, except for two minor reversals. Experiment II also showed that in altering sentences people tended to replace given subjects with subjects higher in the hierarchy by a ratio of about four to one.

There are other important properties of this hierarchy. Consider a Concrete-count noun. When it was substituted for any subject as high or higher in the hierarchy, the resulting sentence was less acceptable than when it was substituted for any subject lower in the hierarchy. According to Experiment I, this was a property of all six types of subject. As a test of this property, 35 comparisons could be made, and all were consistent with this property. Also consider a predicate composed for a Concrete-count subject. It accepted any subject as high or higher on the hierarchy more sensibly than any subject lower on the hierarchy. This property was only slightly less pronounced; out of the 35 possible test comparisons, 32 were consistent with the property. Finally Human subjects could replace each other. Human subjects, in fact, succeeded in this way over all subject types, except the Abstract-count subject, which was based on only 13, not 27 or 40, sampled nouns.

These facts, then, directly support the proposal made in the Introduction: that the logical subjects of transitive verbs form hierarchical categories. This has important consequences for both linguists and psychologists. First, it suggests some important limitations on the selectional restrictions transitive verbs can have. This, in turn, suggests that there is a kind of canonical interpretation of the verb. Second, it suggests a process by which an English speaker might interpret subjects and verbs. And it indicates some limitations on a speaker who is composing a sentence containing a transitive verb. These points will be discussed in turn.

The hierarchy of subject types is of potential interest to linguists, for the hierarchy appears to limit the selectional restrictions of transitive verbs in specific ways. The six noun types studied here can be defined in terms of formal syntactic features, such as [+*Animate*], [−*Count*], and so on. The hierarchy, in fact, can be described quite simply with these features: any noun type with the features, [+*Human*], [+*Animate*], [+*Concrete*], or [+*Count*], is higher in the hierarchy than any type with features containing the corresponding negative value. Human nouns, defined by positive values only on its features, are at the top, and Abstract-mass nouns, with only negative values, are at the bottom.

These facts, along with others that follow, constitute evidence for specific limitations on the verb's selectional restrictions. Whenever a verb allowed [+*Animate*] subjects, it did not restrict them from being [+*Human*]; and whenever a verb allowed [+*Concrete*] subjects, it did not restrict them from being [+*Animate*]. Generally, then, verbs did not restrict their subjects to those having the features [−*Human*], [−*Animate*], or [−*Concrete*]. (The *Count* feature, the most peripheral feature in this study, interacts with *Concrete* according to this generalization.) The generalization is true in the sense that, overall, verbs accepted subjects from a particular point in the hierarchy and up, but did not accept subjects below that point. The limitation, then,

is the following: selectional restrictions can generally only specify that subjects be [+*Human*], [+*Animate*], or [+*Concrete*], and not that they be [−*Human*], [−*Animate*], or [−*Concrete*]. This limitation, of course, is only generally true and would not hold in every particular case.

These limitations suggest a way to measure the "complexity" of a verb's selectional restrictions on its subjects. Lakoff (1965) and Chomsky and Halle (1968) suggest, in phonology, that the two values on a distinctive feature such as voicing are not equal in their complexity value. In the consonants of English, for example, the feature [−*Voiced*] is unmarked, or less complex, and [+*Voiced*] is marked, or more complex. In calculating the complexity of a rule involving consonants, then, the linguist gives less weight to [−*Voiced*] then to [+*Voiced*]. (See also Lakoff, 1966.) The same principle could be applied to the verb's selectional restrictions on subjects. In these restrictions the features [+*Human*], [+*Animate*], and [+*Concrete*], would be unmarked or simple, and their negative counterparts would be marked or complex. This would have the consequence, for example, of differentiating between the relatively "natural" verbs (like *kicked* or *watched*) which are acceptable only to [+*Animate*] subjects, and the relatively "unnatural" verbs (like *subtend* or *coagulate*) which are acceptable only to [−*Animate*] subjects: the former verbs have simpler selectional restrictions than the latter.

These linguistic limitations lead us to a psychological conjecture about the process by which people interpret sentences containing transitive verbs. This process rests on the assumption that the features [+*Human*], [+*Animate*], and [+*Concrete*] are simple, more basic, or more canonical in form than their negative counterparts. The conjecture is this: to intepret any sentence containing a transitive verb a listener implicitly personifies or animates the subject and interprets the sentence as a human or animate subject doing the action. An example shows the possibility of this process. *The map showed the way,* according to this conjecture, is interpretable, because implicitly *map* can momentarily be thought of as human. Since a listener can interpret the sentence with a canonical subject, *The man showed the way,* he is then able to understand that *map* functions like *man,* but with the limitation that *map* is non-human and non-animate. There is a distinct sense in which *map* takes on an animate or human connotation when it is the subject of *show;* this connotation is just as distinctly lacking when *map* is the object of *show,* as in *John showed us the map.*

The conjectured process of interpretation has several advantages. Interpretation is viewed as reducing sentences to simpler underlying forms which can then be understood as canonical forms plus restrictions. This is analogous to the presumed process of understanding superficial sentence strings in terms of base strings plus transformations (see Miller, 1962, and Clark, 1969). The same arguments which support the usefulness of deep structure also support the usefulness of the process of canonical interpretation. Furthermore, the results of the present study follow quite directly from such a process. Sentence subjects, when altered, were changed in the direction of the canonical form; under the present interpretation, this is because they were interpreted in their canonical form. The selectional rules of transitive verbs,

furthermore, were restricted to [+*Human*], [+*Animate*], or [+*Concrete*] subjects—that is, to the canonical values of the features. The process would also explain the grammarian's perception of the subject as actor: the human subject is canonical. Likewise, it would account for the fact that transitive verbs are perceived to be human actions *par excellence*.

Although little is known about how writers compose sentences, the present results suggest some limitations on this process. If composition proceeds, in part at least, from left to right, the choice of a human or animate subject (in an active sentence) allows great freedom in the choice of the verb that follows. On the other hand, the choice of an abstract subject severely limits the choice of the verb that follows; in this case, the verb must be a very general one, one acceptable to all semantic types of subject. Thus a sentence with an abstract subject should generally be much more difficult to generate than one with a human subject. This agrees with the intuitions of the students who composed the sentences in the present study.

Finally, it is important to add one note of caution about the interpretation of the present study. In *Aspects of a Theory of Syntax*, Chomsky (1965) viewed simple sentences with verbs used transitively (like the sentences in the present study) as a homogeneous set. But more recently, other linguists (e.g., Fillmore, 1968; Lyons, 1968) have questioned the homogeneity of this class of sentences. In *The key opened the door*, for example, *the key* is really an "instrumental" used as a subject; *the key* has the same semantic relation in that sentence as it does in *John opened the door with the key*. The "instrumental" relation is semantically quite different from the usual "agentive" relation, as in *John opened the door* (Fillmore, 1968). In the present experiment, however, no distinction was made between the various kinds of subject; all subjects were merely the logical, deep structure subject in Chomsky's sense. On inspection, it was found that, although most subjects were agentive, some, particularly those lower in the hierarchy, were not. In general, this mixture has little effect on the conclusions that have been drawn. It is still true, for example, that Human subjects dominate Animal subjects and that predicates used with hierarchically low subjects are much less specific than those used with high subjects. But, in particular, some results might have two interpretations. For example, Human subjects might dominate Concrete-count subjects partly because the former make better agentive subjects than the latter and partly because sentences with agentive subjects are generally more acceptable than sentences with instrumental subjects. Further research is needed to separate out these potentially important facts.

References

CHOMSKY, N. (1957). *Syntactic Structures* (The Hague).
CHOMSKY, N. (1965). *Aspects of the Theory of Syntax* (Cambridge, Mass.).
CHOMSKY, N. and HALLE, M. (1968). *The Sound Pattern of English* (New York).

CLARK, H. H. (1965). Some structural properties of simple active and passive sentences. *J. verb. Learn. verb. Behav.*, **4**, 365.

CLARK, H. H. (1969). Linguistic processes in deductive reasoning. *Psychol. Rev.*, **76**, 387.

CLARK, H. H. and BEGUN, J. S. (1968). The use of syntax in understanding sentences. *Br. J. Psychol.*, **59**, 219.

FILLMORE, C. J. (1968). The case for case. In E. Bach and R. T. Harms (eds), *Universals in Linguistic Theory* (New York).

JOHNSON, S. C. (1967). Hierarchical clustering schemes. *Psychometrika*, **32**, 241.

LAKOFF, G. (1965). *On the Nature of Syntactic Irregularity.* Computation Laboratory of Harvard University Mathematical Linguistics and Automatic Translation, Report No. NSF-17 (Cambridge, Mass.).

LAKOFF, G. (1966). Stative adjectives and verbs in English. Computation Laboratory of Harvard University Mathematical Linguistics and Automatic Translation, Report No. NSF-17, I-1 to I-16 (Cambridge, Mass.).

LYONS, J. (1968). *Introduction to Theoretical Linguistics* (Cambridge).

MILLER, G. A. (1962). Some psychological studies of grammar. *Amer. Psychol.*, **17**, 748.

SVARTVIK, J. (1966). *On Voice in the English Verb* (The Hague).

THORNDIKE, E. L. and LORGE, I. (1944). *The Teacher's Word Book of 30,000 Words* (New York).