THE NON-UNIQUENESS OF LINGUISTIC INTUITIONS

JOHN M. CARROLL, THOMAS G. BEVER, and CHAVA R. POLLACK
Columbia University

The mental basis of linguistic intuitions is obscure, as regards their relationship both to other aspects of language behavior, such as speaking and listening, and to an hypothesized epistemological structure, such as a 'grammar'. In the present study, we show that experimentally manipulated differences in mental state can systematically alter the linguistic intuitions which speakers render about sentences. These results indicate that the processes underlying intuitions cannot be ignored when they are used as empirical data to test grammatical theories.*

> 'Inner perception ... constitutes the ultimate and indispensable precondition for the other ... sources of knowledge. Consequently ... inner perception constitutes the very foundation upon which the science of psychology is erected' (Brentano [1874] 1973:43) ... 'If anyone were to mount a skeptical attack against this ultimate foundation of cognition, he would find no other foundation upon which to erect an edifice of knowledge. Thus, there is no need to justify our confidence in inner perception. What is clearly needed instead is a theory about the relation between such perception and its object ... such a theory is no longer possible if perception and object are separated into two distinct mental acts, of which the one would only be an effect of the other' (ibid., 140).

INTRODUCTION

1.1. THE NATURE OF INTUITIONS. The most directly available linguistic facts are intuitions about sentences. At least some intuitions can be produced by every native speaker, so they are easily-gathered data in the linguist's own phenomenology. For this reason, theories of language are usually first tested on intuitions about sentences collected from native speaker/hearers of a language. Nevertheless, use of personal intuitions about sentences is not simple at all: the mental basis of such linguistic intuitions is obscure both with respect to how they relate to other aspects of language behavior, such as speaking and listening, and how they relate to an hypothesized epistemological structure, such as a 'grammar'.

Linguists have made the initial assumption that clear intuitions about sentence acceptability and relatedness directly correspond to distinctions made within a grammar, and that such intuitions are related to speaking and listening by way of distinct 'performance' theories which operate on grammatical knowledge. But there is no empirical basis for these assumptions, except the independently verified psychological extensions of the grammar they yield. Un-

* Support was received from a Fellowship from the John Simon Guggenheim Foundation awarded to the second author. Requests for reprints should be directed to the first author, whose present address is 11-105 IBM Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598.

fortunately, behavioral extensions of any currently available grammar beyond the original kind of intuitions which it describes are rare indeed.

In this paper, we consider the implications of the fact that linguistic intuitions are themselves produced via performance mechanisms. Previous discussions have demonstrated this by showing that the linguistic context of an utterance can determine its acceptability (Bever 1970, Carroll 1979, Bolinger 1968, Greenbaum 1977, Labov 1972, Spencer 1973). Furthermore, when intuitions are examined case by case, a bewildering variability appears (Ross 1979). In the present study, we show that experimentally manipulated differences in mental state can systematically alter the linguistic intuitions which speakers render about sentences. The regularity of this effect on intuitions opens up the possibility that the intuitive process can itself be controlled and studied.

Recent achievements in theoretical linguistics have depended primarily on ineluctable intuitions of sentence acceptability and relatedness. For example, it is clear that any syntactic theory must deal with the fact that *John ate the* is not a sentence, and that *John ate the bun* and *The bun was eaten by John* are structurally related sentences. Such intuitions are pre-theoretically solid, and have been properly used as factual touchstones of linguistic theory during the last two decades. During more recent years, however, major schisms have developed in transformational linguistic theory, e.g. between the various schools of 'Autonomous Syntax' on the one hand and of 'Abstract Syntax' on the other.[1] Often, the data cited by each party in such disputes involve intuitions manifestly more subtle than had been used before. If intuitions are themselves in part a function of mental states, it may be that such theoretical disputes are partially related to systematic differences in the intuitions themselves.

In this paper, we present a case study of the commonly-held view that intuitions result from behavioral processes that interact with grammatical structures. Our results suggest that these processes can themselves be systematic, rather than idiosyncratic to each intuition. Accordingly, the study of language requires a psychological theory about its fundamental data as part of the basis for deciding theoretical issues that depend critically on intuitional facts. Our discussion proceeds as follows. First, we apply an established operational technique from social psychology to modify the intuitive process. This technique systematically varies the 'introspective objectivity' which a person applies to a task. We then use this technique to vary the introspective set that speakers apply in judging relatedness among different kinds of causative sen-

[1] We intend 'Abstract Syntax' to refer to the body of syntactic analysis first explored and developed in the works of linguists like George Lakoff, James McCawley, and Paul Postal in the late 1960's and early 1970's (see Postal 1972, Lakoff 1971, for discussion). We intend 'Autonomous Syntax' to refer to the body of syntactic analysis based directly upon Chomsky 1957, 1965, 1977. It is important to note that we are contrasting the 'gist' of these respective schools of linguistic theory, and that we attribute to no single linguist any of the specific analyses we discuss. It is not critical to our argument that the reader accept our characterization of either Abstract Syntax or Autonomous Syntax. We have introduced these labels only in order to focus the discussion. The experimental results we present should certainly NOT be construed as favoring any one position in theoretical linguistics over any other.

tences. Our results show that less introspective objectivity is associated with judgments that are directly consistent with structural descriptions from Autonomous Syntax, while greater introspective objectivity is associated with judgments relatively consistent with structural descriptions from Abstract Syntax.

It should be emphasized that this descriptive result does not impugn or support any particular syntactic theory. It remains an empirical and theoretical question whether intuitions of sentence relatedness are all directly and equally relevant to a syntactic theory. What is important here is not any apparent support of one linguistic theory over another. Rather, it is the fact that a systematic change in the behavioral contexts of intuitions is reflected in a systematic change in the structural descriptions that are most directly implied by those intuitions.

1.2. OBJECTIVE AND SUBJECTIVE MENTAL STATES. Our first task was to find a way to vary systematically a person's introspective state. To do this we drew on the literature concerning the effects of self-observation on human performance. Performance tends to change (usually becoming more intense) when people can observe themselves or others engaged in the same activity. Triplett 1897 experimentally documented this in a study of bicycling: cyclists pedaled faster when they could see other cyclists. Duval & Wicklund 1972 recently integrated the extensive 'social facilitation' literature that followed Triplett's discovery (reviewed by Zajonc 1965) in their theory of 'objective self-awareness'.

The mental state of objective self-awareness, as the name implies, is associated with viewing oneself as a separate object—seeing oneself as others might. The objectively self-aware person is sensitive to a relatively full range of potential situations and interpersonal relations. Conversely, the state of subjective self-awareness is associated with seeing only from inside oneself, with being an unreflective participant in a situation. The subjectively self-aware person is sensitive to a narrower range of situations and relationships. Duval & Wicklund present a number of experimental studies which clarify and operationalize the distinction between these states.

It is striking how similar the state of 'objective' self-awareness is to the mental state of the linguist while rendering intuitions. Speaker/hearers are subjectively preoccupied; they are not ordinarily aware of unacceptabilities, presuppositions, and ambiguities in their own utterances, or in the speech of others. In order to have intuitions about acceptability, linguists must almost cease to be behaving speaker/hearers. They must pause and reflect; in the terms of Duval & Wicklund, linguists objectify the sentence from all the specific potential functional contexts of its utterance. Duval & Wicklund themselves suggest (p. 6) that the objectively self-aware person might be especially sensitive to grammatical correctness. But what exactly does 'correctness' mean in linguistic enquiry?

The question highlights an important difference between the linguist's intuitions and the behaviors studied by Duval & Wicklund. In their experimental

tasks, correctness was always defined objectively (e.g. in terms of the accuracy or speed of performances). This made it possible for them to measure quantitatively the behavioral effects of changes in objective self-awareness. But no corresponding a-priori definition exists for a 'correct' intuition about sentences. While we cannot measure quantitative changes in the 'correctness' of linguistic intuitions, we may be able to identify qualitative changes in intuitions resulting from changes in objective self-awareness. Having intuitions at all requires a minimum level of objective self-awareness. However, persons who are relatively more objectively preoccupied might experience different intuitions from those who are relatively less objectively preoccupied. In particular, a more objectively preoccupied speaker may be more sensitive to generalizations based on interpersonal linguistic experience.

1.3. THE MIRROR MANIPULATION. In testing this prediction for linguistic intuitions, our first problem was to decide on a technique that would manipulate objective mental state, without our consultants' being aware of it. Duval & Wicklund's standard technique to elicit the mental state of objective self-awareness is to place a mirror before an experimental subject. Subjects are not informed that the mirror has any significance; however, the mere fact that they can see themselves (as objects) reflected in the mirror facilitates a mental state of objective self-awareness; this is often reflected in performance changes on an experimental task.

We initially explored the effect of objective self-awareness on linguistic intuitions other than those of acceptability or similarity. In our first study, we asked subjects to rate the 'truth' of statements like these:

(1) a. A house is a building.
    b. A garage is a building.
    c. A lean-to is a building.
    d. A tent is a building.
(2) a. A robin is a bird.
    b. A chicken is a bird.
    c. An ostrich is a bird.
    d. A bat is a bird.

In these sentence sets, the subject nouns differed systematically as to how 'focally' they instantiate the category denoted by the object noun (Rosch 1975).

Each consultant rendered a judgment for one of the sentences in eight sentence sets like the examples above. Consultants were given booklets and allowed to work at their own speed, alone in small experimental rooms, assigning truth judgments on a 1 to 10 scale to each sentence pair. Each time a consultant decided on a judgment for a particular pair, he or she had to consult a 'key' which indicated where on the answer sheet to enter the judgment.

Ten of the consultants performed this rating task while seated in front of the mirror, to which the 'key' was taped. Ten of them performed the task with the mirror turned around and the 'key' taped to its back. In each group, five of the subjects were male, and five were female. All were native speakers of

English between the ages of 18 and 25 years, with no reported training in linguistics.

We expected that the truth judgments would vary as a function of the degree to which the subject noun was an archetypical member of the category named by the predicate-object noun. In addition, we predicted that speakers in the objective self-awareness (or OSA) condition would rate all the sentences as 'truer' than the speakers in the subjective self-awareness (or SSA) condition. This follows from the fact that each sentence is 'true' in SOME communicative situation, and the more objectively self-aware person will tend to consider more such situations. The mean ratings (averaged across speakers and sentence sets) for the two conditions are presented in Table 1 ('10' is most true, '1' is least true).

SENTENCE TYPE	SSA CONDITION	OSA CONDITION
A house is a building.	9.55	9.60
A garage is a building.	7.60	8.60
A lean-to is a building.	5.40	7.35
A tent is a building.	3.25	4.50
OVER-ALL	6.45	7.51

TABLE 1.

The ratings of truth in the OSA condition were reliably higher than those in the SSA condition.[2] Further, the OSA/SSA difference was larger for the (c) sentences; this suggests that the OSA condition consultants distinguished more between the archetypically unlikely but technically true (c) sentences and the technically false (d) sentences than did the SSA condition consultants. These results confirmed our predictions and lent initial validity to the mirror technique as a means of manipulating introspective set vis-à-vis linguistic intuitions.

Before we turn to an application of this technique to intuitions of sentence relatedness, it is useful to consider the implications of these results for semantic investigations. Thus Lakoff 1972 argues that data like Rosch's and those in Table 1 demonstrate that 'truth' is not a categorical concept, but rather a 'fuzzy' gradation of 'more and less true'. In the light of our findings, should we now conclude that the formal characterization of linguistic truth must also refer to the presence of mirrors—or, at least, to OSA 'levels'? Obviously not: this determination remains a theoretical question. Our results extend the empirical basis for theoretical analyses of linguistic truth, and set constraints on what could be a comprehensive account; but they do not dictate their own description in linguistic theory. Indeed, we found a large and highly significant difference between truth judgments rendered for sentences (a)–(c), on the one hand, and those rendered for the (d) sentences, on the other, in BOTH the OSA

[2] Consultants in the OSA condition rendered higher average truth ratings than informants in the SSA condition ($p < .05$, by Wilcoxon's Signed-Ranks Test for Uncorrelated Samples). Sentence sets also received higher average truth ratings in the OSA condition ($p < .025$, by Wilcoxon's Matched-Pairs Signed-Ranks Test). (See Ferguson 1971 for discussion of Wilcoxon's non-parametric statistical tests.)

and the SSA conditions.[3] Prima facie, this latter result reconfirms the possibility that categorical distinctions of truth and falsity CAN exist in behavior.

INTUITIONS ABOUT INTERSENTENTIAL RELATIONS

2. Intuitions about intersentential relations are important data in the discovery and verification of grammars. Chomsky (1965:21–4) defines a 'descriptively adequate' grammar as one that provides an account of such intuitions. He gives as examples the relations of ambiguity and paraphrase; but these relations do not exhaust the criterion of descriptive adequacy. A grammar is justified to the extent that it accurately describes ALL intersentential relations that are salient in the linguistic intuitions of speaker/hearers. Recent linguistic research refers to an increasing variety of such intersentential relations (e.g. entailment, presupposition, and cohesion).

Psycholinguistic research has also evaluated grammars in terms of their ability to account for intuitions about intersentential relations such as 'similarity' between sentences. In the first major study of this sort, Clifton & Odom 1966 asked subjects to rate the similarity of pairs of structurally related sentences. The pairs were drawn from sets of eight sentences. Each set consisted of a simple transitive sentence with its passive, its negation, its question, and the combinations of these (e.g. negative-passive-question). For each set, one member was designated as the 'standard', against which the remaining seven sentences were rated for similarity. Clifton & Odom used the patterns of similarity judgments to compare predictions based on two syntactic models: Chomsky 1957 vs. Katz & Postal 1964. These models gave different structural analyses for the eight types of constructions.

Clifton & Odom subjected means of the 28 pair-wise similarity judgments to a multidimensional scaling analysis. (For discussion of multidimensional scaling techniques, see Kruskal 1964, as well as the Appendix below.) Multidimensional scaling transforms a matrix of numerical relationships—in this case, the similarity judgments rendered by their subjects—into a plot of spatial relationships. The plot of Clifton & Odom's data on sentence similarity showed that all the sentence forms were roughly equidistant, with the exception of two pairs which were virtually coincident: the question and negative-question, and the passive-question and passive-negative-question.

On the analysis presented in Chomsky 1957, this result has no obvious explanation. In that analysis, the negative, question, and passive constructions (and their combinations) were derived by optional transformations, with the same underlying structure in each case. There are two obvious ways to interpret how such a theory will be reflected in a spatial representation of similarity ratings: similarity is based either on transformational identity, or on identity of underlying structure. The former would predict that the sentence forms lie at the corners of a 3-dimensional 'cube', with each dimension corresponding to one of the transformations (see Miller 1962). Alternatively, if 'similarity' is

[3] For speakers in both conditions, $p < .001$, using the Sign Test. (See Ferguson for discussion of the Sign Test.)

based on underlying structural identity, then the sentence forms should be mutually equidistant. Neither interpretation is consistent with the results obtained, since all the pairs but two were equidistant.

The view that similarity ratings are based on underlying structure is directly consonant with the Katz & Postal analysis of these sentence forms. In that analysis, the underlying structures for question and negative-question and for the passive-question and negative-passive-question are identical, while the other forms have different underlying structures from each other. Clifton & Odom concluded that the sentence-similarity data supported Katz & Postal's analysis over Chomsky's.

One of Clifton & Odom's principal concerns was to distinguish between two syntactic theories. This purpose may well have been outmoded by developments in linguistic theory. We will use their experimental task for another purpose entirely: namely, to study the patterns of intuitions about sentence relations that occur under different conditions. On our view, speakers adopt particular strategies for interpreting their intuitions, depending on immediate behavioral context. These strategies emphasize particular aspects of linguistic structure, which in turn form the basis for personal similarity judgments. We concerned ourselves with patterns of relations among sentences like these:

(3) a. ACTIVE: The morning sun dried the sweet raisins.
b. PASSIVE: The sweet raisins were dried by the morning sun.
c. INCHOATE: The sweet raisins dried in the morning sun.
d. WERE-INCHOATE: The sweet raisins were dried in the morning sun.
e. CAUSE: The morning sun caused the sweet raisins to dry.
f. BECAUSE: The sweet raisins dried because of the morning sun.

(The labels we have placed on these sentences are for convenience only; they should not be interpreted as implying formal analyses of the structures.)

Of course, many descriptions of the relations among these sentences are possible. In this paper, we will use the different descriptions offered by two schools of transformational grammar: Autonomous Syntax and Abstract Syntax. There are many differences between these two schools, and many variants within each school. The salient descriptive feature for our concern here involves the underlying source for the sentences. Followers of Abstract Syntax would argue that many of these six sentence types derive from identical (or very similar) underlying structures, because of their semantic similarity. For example, on this analysis, the actual CAUSE sentence is close in form to the structure that underlies both the ACTIVE and BECAUSE sentences, as well as the CAUSE sentence itself (see below).[4] In contrast, followers of Autonomous Syn-

[4] At first, it might seem that an Abstract Syntax analysis would assign INCHOATE and WERE-INCHOATE the same underlying structure as the other four sentences, just because of their similarity in meaning. Closer inspection, however, shows that this analysis would be incorrect. Both the INCHOATE and WERE-INCHOATE contain an unspecified agent other than *the sun*—as exemplified in the sentences below, into which optional agents have been inserted:

The sweet raisins dried in the morning sun because of the wind.
The sweet raisins were dried by the wind in the morning sun.

tax would analyse these three sentence types as deriving from different deep structures.[5]

These two theories are intended to be structural accounts of sentences. However, we can differentiate the theories according to the intuitive linguistic level they emphasize. The level which Abstract Syntax analysis addresses is, by design, more communicatively generalized than that which Autonomous Syntax addresses. For example, analyses within Abstract Syntax represent structurally the implicit causation expressed in a sentence like 3a, while analyses in Autonomous Syntax represent only the exchange of major NP's that relates forms like 3a–b. We decided to examine the bases of these factual disputes by referring to the OSA/SSA manipulation. It appeared that the pattern of intuitions supporting each kind of analysis might reflect a systematic difference in a speaker's current self-awareness.

We postulated that the OSA state would encourage an Abstract Syntax analysis as a strategy for organizing intuitions; and that a state of relatively less objective self-awareness would encourage a strategy of interpreting intuitions according to an Autonomous Syntax analysis. This differentiation follows from the hypothesis that objectively self-aware speakers have an implicitly

---

The other four sentence types do not have an agent missing, as demonstrated by the unacceptability of the following sentences:

*The morning sun dried the sweet raisins because of the wind.
*The sweet raisins were dried by the morning sun by the wind.
*The morning sun caused the sweet raisins to dry because of the wind.
*The sweet raisins dried because of the morning sun because of the wind.

In none of these sentences is *wind* interpretable as the causative agent of the activity of *drying*. Accordingly, the INCHOATE and WERE-INCHOATE sentences differ from the other four in their underlying relations. They also differ from each other. The WERE-INCHOATE requires an agent in its underlying form, while the INCHOATE does not; i.e., the underlying source for the WERE-INCHOATE would be something like *NP caused (the raisins become dry) in the sun*, while the underlying structure for the INCHOATE would be like *The raisins became dry in the sun.*

One might want to argue that an Abstract Syntax analysis would also assign a reading to the INCHOATE and WERE-INCHOATE that derives from an underlying structure identical to the other four sentence types we are studying. Even if this were so, it would not change our predictions about the scaling analysis of similarity. Scaling is sensitive to RELATIVE distances among stimuli. Since the INCHOATE and WERE-INCHOATE are different from the other four types on at least one reading, they should be scaled as distant from those types.

[5] Cf. fn. 1, above. It should be noted that the Autonomous Syntax theory went through a phase, the so-called Standard Theory (Katz & Postal 1964, Chomsky 1965), in which ACTIVE, PASSIVE, CAUSE, and BECAUSE sentences would all be viewed as sharing a common source. One reason for this is that, to a first-order approximation, the sentences are semantically identical: they all propose *the sun* as the agent of the *raisins* becoming *dry*. In Standard Theory, such semantic similarity was directly explained by reference to a common deep structure, since deep structures were the sole source of relational semantic information. Later versions of Autonomous Syntax (e.g. Extended Standard Theory, Chomsky 1971; Trace Theory, Chomsky 1977, Chomsky and Lasnik 1977) have been able to return to a deep-structure analysis similar to that of Chomsky 1957, with respect to the issues under discussion here, because identity of semantic relations no longer presupposes identity of deep structure. (But see Katz 1972 for another solution, maintaining the sole semantic role of deep structures but allowing the same deep-structure differences as in Chomsky 1957.)

heightened sense of social interaction, and therefore will be more sensitive to the communicative similarity among the sentences.

2.1. QUANTITATIVE DIFFERENCES IN SIMILARITY RATINGS. There are 15 possible pairings of the six sentence types listed above (e.g. ACTIVE with PASSIVE and CAUSE with ACTIVE). We constructed 15 such sentence sets, each with a different actor, action, and object. Consultants were asked to render similarity judgments for the 15 possible pairings of sentence types, each of which was drawn from a different sentence set. Thus there were 15 different presentations of 15 sentence pairs each; each presentation included one sentence pairing from each sentence set, and one instance of each of the possible pairings of sentence types.

The presentations were arranged in a Latin square with respect to sentence-type pairings, so that each possible pairing appeared once in each of the 15 serial positions in a presentation. The 15 presentations were prepared as booklets, with a different pair appearing on each page. Two of 30 consultants were randomly assigned to each of the 15 presentations. One consultant in each of the 15 presentation conditions was assigned to the OSA condition, and the other was assigned to the SSA condition. Again, participants were native speakers of English with no reported training in linguistics. Both the OSA and SSA groups consisted of eight males and seven females.

The predicted effect of OSA did in fact occur: consultants rated the similarity of sentence pairs as greater when in the OSA condition than in the SSA condition (see Table 2). We tested this effect statistically by speaker, by sentence

	SSA CONDITION	OSA CONDITION
ACTIVE VS. PASSIVE	8.66	7.80
ACTIVE VS. INCHOATE	6.66	6.73
ACTIVE VS. WERE-INCHOATE	5.86	6.60
ACTIVE VS. CAUSE	7.26	7.46
ACTIVE VS. BECAUSE	6.60	7.53
PASSIVE VS. INCHOATE	6.93	6.86
PASSIVE VS. WERE-INCHOATE	6.80	7.46
PASSIVE VS. CAUSE	5.93	7.80
PASSIVE VS. BECAUSE	7.06	7.73
INCHOATE VS. WERE-INCHOATE	8.13	6.86
INCHOATE VS. CAUSE	5.73	6.00
INCHOATE VS. BECAUSE	5.26	6.73
WERE-INCHOATE VS. CAUSE	6.06	6.80
WERE-INCHOATE VS. BECAUSE	5.40	6.40
CAUSE VS. BECAUSE	7.80	8.06
OVER-ALL	6.76	7.12

TABLE 2.

set, and by sentence-type pairing.[6] As in our first experiment, this indicates that linguistic intuitions can systematically vary as a function of the mental

[6] Consultants in the OSA condition assigned higher average ratings (mean = 7.12 out of 10.0) than those in the SSA condition (mean = 6.76 out of 10.0; $p < .05$, by Wilcoxon's Matched-Pairs Signed-Ranks Test, matching subjects in the same presentation condition). Sentence sets received

state of OSA. In particular, the results of this experiment suggest that the more objectively preoccupied person is more sensitive to the communicative similarity between sentences, and accordingly assigns higher similarity ratings to sentences that are communicatively similar. The less objectively preoccupied person is less sensitive to the relations, and assigns lower similarity ratings.

2.2. QUALITATIVE DIFFERENCES IN SIMILARITY RATINGS. We next considered the qualitative differences between the OSA and SSA conditions. In order to do this, we performed a multidimensional scaling analysis on the similarity judgments. (See our discussion of Clifton & Odom, above.) The scaling plot is 'non-metric'; i.e., the distances are linearly related not to the actual values in the original matrix, but rather to the ordering relationships between these values. In the present case, this means that multidimensional scaling plots of the OSA condition and the SSA condition judgments can be compared directly, in spite of the higher over-all similarity judgments elicited by the OSA condition. The scaling plots represent the 'qualitative' properties of the judgments obtained in the two conditions, REGARDLESS OF THE OVER-ALL QUANTITATIVE DIFFERENCE.

It is useful to know ahead of time what spatial organization can be expected in each condition. We have used the over-all difference in similarity ratings to suggest that the SSA condition brings out an organization of intuitions associated with Autonomous Syntax, and the OSA condition one with Abstract Syntax. We follow Clifton & Odom in assuming that speakers assign similarity ratings according to their current interpretation of similarity in underlying structure. This predicts that the ACTIVE and PASSIVE will be judged highly similar, and different from the other constructions in the SSA condition: on an Autonomous Syntax analysis, only the active and passive constructions share a deep structure, among the six constructions we are using.

Figure 1 (overleaf) presents a mean two-dimensional plot of the similarity data for the SSA condition. (Concerning our use of multidimensional scaling, see also the Appendix, below. In the text we only discuss the data as averaged across several analyses in each condition.) This representation has two striking features. First, the six sentence types appear to cluster into one strong group and two weak ones; and these three groups are mutually about as far apart as they can be. Second, within these three groups, PASSIVE and ACTIVE are much nearer one another than are CAUSE and BECAUSE, or INCHOATE and WERE-INCHOATE; this pattern is highly significant.

If judgments are based on similarity in underlying structure, then an ACTIVE–PASSIVE cluster is exactly that predicted by an Autonomous Syntax analysis, in which only those two sentence forms share an underlying structure. Although the other two 'clusters' are not statistically reliable, they may indicate

---

higher average rating in the OSA condition than in the SSA condition ($p < .005$, by Wilcoxon's Matched-Pairs Signed-Ranks Test). Finally, sentence-type pairings received higher ratings in the OSA condition than in the SSA condition ($p < .025$, by Wilcoxon's Matched-Pairs Signed-Ranks Test).

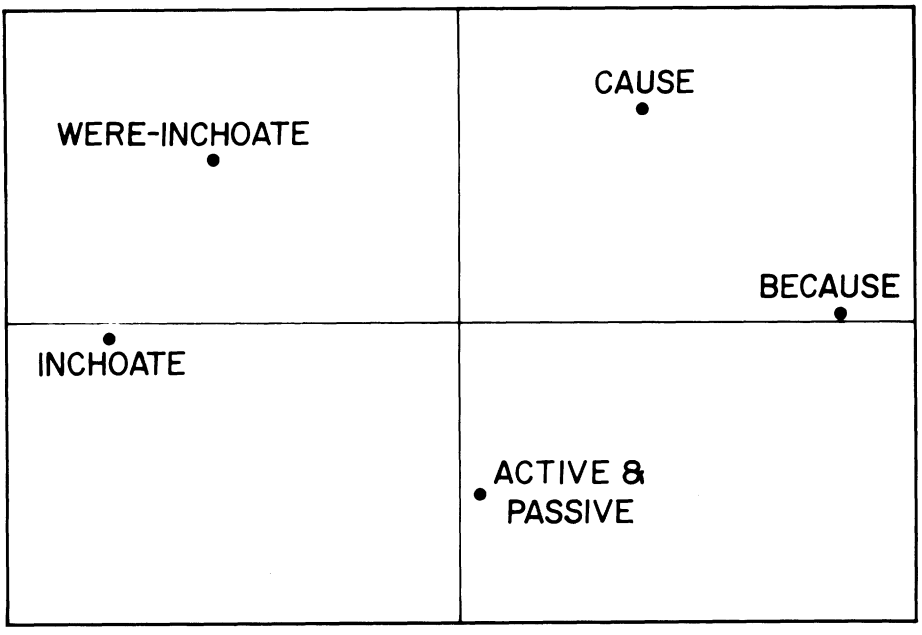

FIGURE 1. Each point is the mean position of 11 low-stress plots.

that surface identity of morphemes is also part of the over-all criterion of similarity. INCHOATE and WERE-INCHOATE differ by only one word (namely, *were*) in their lexical manifestations. Further, CAUSE and BECAUSE are the only sentence types that explicitly have the root *cause*.

The predictions made for the OSA condition by an Abstract Syntax analysis are more complicated. First, consider the group CAUSE, BECAUSE, ACTIVE, and PASSIVE. In an Abstract Syntax framework, all these sentence types can be derived from a single underlying semantic structure, bracketed as follows:

(4) (CAUSE SUN (DRY RAISINS))

If no optional transformations apply to this structure, it is realized by a surface sequence like the CAUSE sentence. If the predicate DRY is raised into the main clause, it will be realized by a surface sequence like the ACTIVE sentence (and from this structure, the PASSIVE sentence would be derived as in Autonomous Syntax). Finally, if the entire embedded clause (DRY RAISINS) is raised out of the bracketing structure, itself becoming the main clause, the surface sequence realized will be like the BECAUSE sentence.

The INCHOATE and WERE-INCHOATE sentences have quite different analyses (see fn. 4). The INCHOATE form has an intransitive verb, while the sentence types discussed above have either a transitive verb or an intransitive verb plus an explicit use of *cause* (or *because*). The WERE-INCHOATE form is also quite different from the previously-discussed sentence types. On the most obvious reading, it is a truncated passive:

(5) The sweet raisins were dried in the morning sun BY SOMEONE/SOMETHING.

Hence, though it contains a transitive form of the verb *dry*, it has a different agent than the four sentence types which have *the sun* as the agent. These considerations predict that, in the OSA condition, ACTIVE, PASSIVE, CAUSE, and BECAUSE will be judged as identical to each other, while INCHOATE and WERE-INCHOATE will each be judged different from all the other constructions.

Figure 2 presents the mean of the reliable two-dimensional scaling plots for the judgments obtained from the OSA condition subjects. This plot is exactly that predicted by underlying identity of structure in an Autonomous Syntax analysis: ACTIVE, PASSIVE, BECAUSE, and CAUSE cluster together, while INCHOATE and WERE-INCHOATE are isolated points.

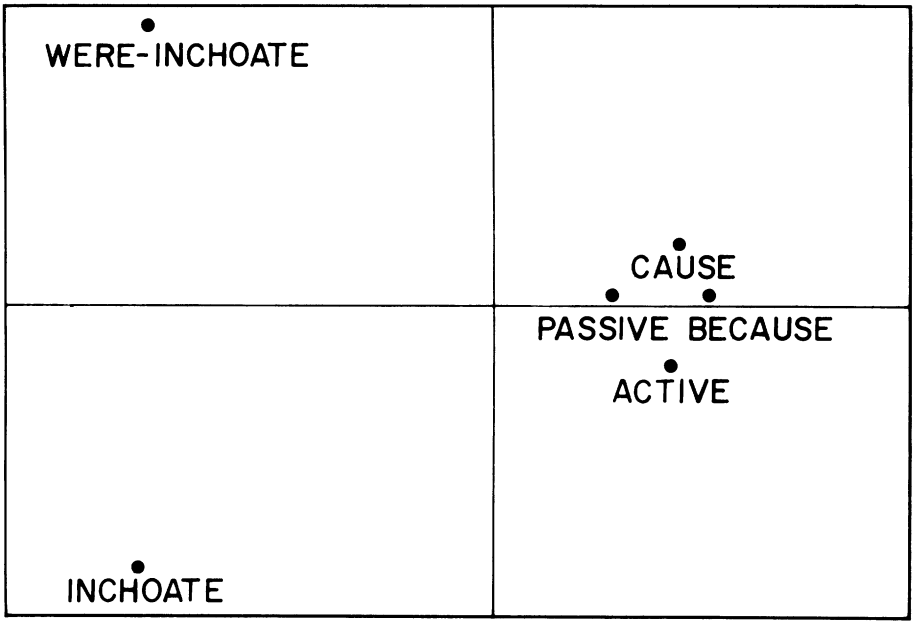

FIGURE 2. Each point is the mean position of 15 low-stress plots.

We conclude from these scaling analyses that the qualitative differences between the OSA and SSA conditions correspond to differences between the Abstract and Autonomous Syntactic analyses of these sentences.

IMPLICATIONS

3. Systematically different sets of intuitions about sentence relatedness can exist in the same linguistic population. This raises the possibility that there are relatively stable mental schemata that underlie manifest intuitions of similarity. On this view, speakers shift their intuitional processes from one form to another without being aware of doing so. We do not have enough cases to motivate a full theory of how intuitional processes interact with linguistic structure. However, we can start the investigation of such a theory by postulating that speakers use INTUITIONAL STRATEGIES to organize their introspective linguistic percepts. Consultants who are objectively self-aware adopt a communicative strategy,

which emphasizes semantic similarity. Those who are in an SSA state use a sentential strategy, relying on more focused structural properties of sentences.

It is startling that we can identify intuitions rendered under the communicative strategy with a particular linguistic analysis, and those rendered under the sentential strategy with a different linguistic analysis of the same data. This suggests that some of the apparent empirical support for factual conflicts between these analyses could reflect a covert shift in intuitional strategy; i.e., the analyses could both be correct because they are, in fact, analyses of different linguistic phenomena (see Katz & Bever 1976 for fuller discussion of this).

Accordingly, the interaction between linguistic intuitions and the OSA state does not selectively support any particular school of linguistic theory. What these results do indicate is that the process of having linguistic intuitions cannot be ignored in their use as empirical data to test theories. Our experiment should be taken as an emphatic demonstration that linguistic intuitions have a dual systematic nature. On the one hand, they can be basic and primitive manifestations of the grammatical knowledge speakers share; but on the other hand, they are complex behavioral performances that can be properly understood and adequately interpreted only by a comprehensive analysis.

Many questions remain unanswered concerning the behavioral basis of rendering intuitions. Psychologists still fear the excesses of Titchenerian introspectionism (Boring 1953); and, by common consent, they have outlawed intuition as an empirical source, rather than trying to understand this type of behavior. A full explication of the present results must include a theory of how human knowledge and behavior interact to configure manifest intuitions. Trying to avoid the problem by pre-theoretic fiat will not remove the effects described by such a theory; it will only make them unpredictable.

Some may feel our data show that the empirical apocalypse is at hand for linguistics: if intuitions are not the factual bedrock of linguistic science, then we can ask, is it a science of anything at all? Fortunately, such a quandary is unwarranted by our data. We have shown that relatively subtle intuitions can be systematically manipulated. But the science of linguistics is the science of LANGUAGE, not merely the orderly description of intuitions. The pre-theoretically clear intuitions such as those cited in our introduction remain valid facts, which any description of language must include. Furthermore, many linguistic structures may be discovered deductively, on theory-internal grounds, rather than based on intuitional data.

If we look both backward and forward, this is the state of the art, as we understand it. Solid, pre-theoretic sentence intuitions were used as the factual test for initial theoretical developments in generative grammar. As theories have become more sophisticated, the facts they aspire to describe have become more subtle. However, the expression of these subtle, but theoretically crucial, aspects of linguistic knowledge can be systematically distorted in the process of rendering linguistic intuitions. This both makes possible and requires the use of a theory of introspection as part of the theoretical interpretation of

intuitions. Such a theory will articulate a richer understanding of language behavior—embracing listening, talking, AND THE INTUITIONAL PROCESS ITSELF.

APPENDIX

A word is in order about our use of multidimensional scaling, in particular with reference to the statistical significance of the findings. There are two questions: First, is each pattern reliable for the solution in two dimensions? Second, are the two cluster patterns significantly different?

Multidimensional scaling is non-deterministic: the same similarity data on 15 pairs of sentences can result in different solutions, depending on the order in which the similarity values are accessed. Each solution has a different goodness-of-fit to an ideal representation: this is indicated in an index for each plot called 'stress', a least-squares measure of the amount of variability in the data for which the solution does not account. Hence, the lower the stress, the less the solution violates the ordinal relations among the original similarity ratings.

Often the goal is to prove that a particular set of similarity data is representable in a particular number of dimensions. For example, we might wish to demonstrate simply that the six sentences CAN be plotted in an n-dimensional analysis that accounts for a significant amount of the variance. To do this, we simply must show that the stress level assigned to an n-dimensional plot is significantly lower than chance, and that the stress level is higher at $n + 1$ dimensions and not appreciably lower at $n - 1$ dimensions. This situation does obtain for the 2-dimensional analyses of both the OSA and SSA conditions: the stress values we obtained for sample 1-dimensional solutions were .285 for the SSA condition and .375 for the OSA condition. Since our program automatically terminated when stress reached .01, and since most of our 2-dimensional solutions attained stress values this low, stress values of sample 3-dimensional solutions were not lower than our 2-dimensional results. This gives initial validity to the claim that the data can be appropriately represented in two dimensions.

However, we must be cautious in concluding this. A problem arises in multidimensional analysis when the number of dimensions approaches the number of points being plotted: obviously, a 6-dimensional plot can always account for six data points with zero stress, even if they actually have no systematic underlying organization. Even a 2-dimensional plot of six actually random datapoints can occur with a low stress some of the time. The frequency of such spurious low-stress solutions has been estimated by Monte Carlo simulation. Klahr 1969 (see also Wagenaar & Padmos 1971) showed that six actually unstructured points will appear to have a 2-dimensional solution with a final stress $\leq .01$ on 20% of the analyses.

We tested the 2-dimensional reliability of our solutions for OSA and SSA conditions in the following way. Each set of data was analysed 18 different times with a randomly different starting configuration, and thus access order, of the 15 similarity ratings. To be conservative, we did not use the frequently-used TORSCA configuration that avoids local minima. Our program was KYST (Kruskal, Young, Shepard, Torgerson), a package at the Columbia University Center for Computing Activities, which is basically a merging of M-D-SCAL 5M and TORSCA 9; it automatically terminates each solution after (a) a solution with stress $\leq .01$ is found, (b) a repeating stress level occurs 10 times, or (c) 50 passes through the matrix have failed to produce (a) or (b).

We obtained 15 solutions with stress $\leq .01$ for the OSA data, and 11 solutions for the SSA data. This contrasts with the mean of 3.6 solutions of stress $\leq .01$ expected if the input data were random. (Chi square = 42.02 and 16.53, df = 1, $p < .001$ for both conditions.) That is, although the number of points studied is small, the fact that there is a high goodness-of-fit of the solutions shows that the data in each condition are meaningfully represented in two dimensions.

But how reliable is the particular pattern in the plot for each condition, and how reliably does each pattern conform to its hypothetically corresponding grammatical analysis? Figs. 1–2 above actually present the composite mean position for each of the six points in the plots with stress $\leq .01$ (reflected for symmetrical solutions when relevant). Table A (overleaf) lists the standard deviations of the direct line distances from each of those means; the low S.D.'s suggest that the pattern is reliable.

It is possible to determine HOW reliable each pattern is. The data underlying Table A allow a

	SSA			OSA		
COMPARISON	MEAN	S.D.	Z	MEAN	S.D.	Z
ACTIV/PASSV	.046	.063	.73	.314	.281	1.12
ACTIV/CAUSE	1.446	.143	10.11	.532	.451	1.18
ACTIV/BECAU	1.438	.140	10.27	.282	.285	.99
PASSV/CAUSE	1.440	.112	12.85	.289	.240	1.20
PASSV/BECAU	1.459	.146	9.99	.312	.275	1.35
CAUSE/BECAU	1.014	.441	2.30	.354	.338	1.05
ACTIV/INCHO	1.446	.143	10.11	1.923	.139	13.83
ACTIV/WERIN	1.441	.134	10.75	2.079	.081	25.67
PASSV/INCHO	1.423	.145	9.81	1.824	.197	9.26
PASSV/WERIN	1.450	.136	10.66	1.797	.192	9.36
CAUSE/INCHO	1.970	.295	6.68	2.098	.075	27.97
CAUSE/WERIN	1.468	.133	11.04	1.922	.113	17.00
BECAU/INCHO	2.452	.238	10.30	2.148	.081	26.52
BECAU/WERIN	2.229	.204	10.73	2.081	.052	40.02
INCHO/WERIN	.695	.548	1.27	1.815	.172	10.55

TABLE A. 'Z' in S.D.'s refers to the distance from zero between the members of a pair.

quantification of the strength of the clusters. In the SSA condition, all 11 plots show ACTIVE and PASSIVE closer than any other pair ($p < .0005$ by a Sign Test); INCHOATE and WERE-INCHOATE are closer to each other than to CAUSE or BECAUSE ($p < .0005$ for each), and CAUSE and BECAUSE are closer to each other than to INCHOATE or WERE-INCHOATE ($p < .003$). INCHOATE and WERE-INCHOATE, and CAUSE and BECAUSE sentences show only a tendency ($.033 < p < .327$) to be closer to one another than to ACTIVE or PASSIVE. That is, the ACTIVE–PASSIVE cluster is statistically reliable, while the INCHOATE–WERE-INCHOATE and CAUSE–BECAUSE clusters are not.

In the OSA condition plots, both INCHOATE and WERE-INCHOATE are always further from each of the other four points than any pair of the other points are from each other ($p < .00005$ for each comparison). As we have pointed out, the plots are of relative, not absolute, position; hence the significant distance of INCHOATE and WERE-INCHOATE from the other four points, and from each other, reflects the fact that the other four points themselves form a significant cluster.

The data in Table A offer further support for these cluster analyses. We can quantify the possibility that the members of a cluster actually lie on the same point by calculating the z-score for distance = zero, relative to the mean distance between that pair. For example ACTIVE and CAUSE in the SSA condition have a z-score for zero distance of 10.11, meaning that their actual distance apart is 10.11 standard deviation units away from zero—a highly significant distance. In contrast, the mean z-score for zero distance of the 'cluster' ACTIVE and PASSIVE in the SSA condition is .73, a non-significant difference. In the OSA condition, all six 'cluster' pairs from among ACTIVE, PASSIVE, CAUSE, and BECAUSE have a z-score for distance = zero ≤ 1.35; the 'non-cluster' pairs have a corresponding score ≥ 9.26. That is, the members of the clusters are not significantly different from lying on the SAME point, while non-clustering sentence forms are significantly far apart.

In brief, careful consideration shows that our analyses are reliable; that the two conditions differ significantly; and that they each conform to a distinct theoretical prediction based on a distinct linguistic analysis. The reason we did not use more sentence types to avoid statistical problems is both practical and theoretical. For example, 10 sentence types would have required 45 separate sentence pairs for judgment—a considerable strain for both experimenter and subject. Even worse, the effect of making many judgments could in itself be expected to heighten the subjects' OSA level, perhaps washing out the effect of the mirror manipulation.

REFERENCES

BEVER, THOMAS G. 1970. The cognitive basis for linguistic structures. Cognition and the development of language, ed. by John R. Hayes, 279–362. New York: Wiley.
BOLINGER, DWIGHT. 1968. Judgments of grammaticality. Lingua 21.34–40.

BORING, EDWIN G. 1953. A history of introspection. Psychological Bulletin 50.169–89.
BRENTANO, FRANZ. 1973. Psychology from an empirical standpoint. New York: Humanities Press. [Originally published as Psychologie vom empirischen Standpunkt. Leipzig: Duncker & Humblot, 1874.]
CARROLL, JOHN M. 1979. Complex compounds: Phrasal embedding in lexical structures. Linguistics 17.863–77.
CHOMSKY, NOAM. 1957. Syntactic structures. The Hague: Mouton.
——. 1965. Aspects of the theory of syntax. Cambridge, MA: MIT Press.
——. 1971. Deep structure, surface structure, and semantic interpretation. Semantics, ed. by Danny D. Steinberg & Leon A. Jakobovits, 183–216. Cambridge: University Press.
——. 1977. Essays on form and interpretation. New York: American Elsevier.
——, and HOWARD LASNIK. 1977. Filters and control. Linguistic Inquiry 8.425–504.
CLIFTON, CHARLES, and RICHARD ODOM. 1966. Similarity relations among certain English sentence constructions. Psychological Monographs 80, whole no. 613.
DUVAL, SHELLEY, and ROBERT A. WICKLUND. 1972. A theory of objective self awareness. New York: Academic Press.
FERGUSON, GEORGE ANDREW. 1971. Statistical analysis in education and psychology. New York: McGraw-Hill.
GREENBAUM, SIDNEY. 1977. Contextual influence on acceptability judgments. Linguistics 187.5–11.
KATZ, JERROLD J. 1972. Semantic theory. New York: Harper & Row.
——, and THOMAS G. BEVER. 1976. The fall and rise of empiricism. An integrated theory of linguistic ability, ed. by Thomas Bever et al., 11–64. New York: Crowell.
——, and PAUL M. POSTAL. 1964. An integrated theory of linguistic descriptions. Cambridge, MA: MIT Press.
KLAHR, DAVID. 1969. A Monte Carlo investigation of the statistical significance of Kruskal's non-metric scaling procedure. Psychometrika 34.319–30.
KRUSKAL, JOSEPH B. 1964. Multidimensional scaling by optimizing goodness of fit to a non-metric hypothesis. Psychometrika 29.1–27.
LABOV, WILLIAM. 1972. Sociolinguistic patterns. Philadelphia: University of Pennsylvania Press.
LAKOFF, GEORGE. 1971. On generative semantics. Semantics, ed. by Danny D. Steinberg & Leon A. Jakobovits, 232–96. New York: Cambridge University Press.
——. 1972. Hedges: A study in meaning criteria and the logic of fuzzy concepts. CLS 8.183–228.
MILLER, GEORGE. 1962. Some psychological studies of grammar. American Psychologist 17.748–62.
POSTAL, PAUL M. 1972. The best theory. Goals of linguistic theory, ed. by Stanley Peters, 131–70. Englewood Cliffs, NJ: Prentice-Hall.
ROSCH, ELEANOR. 1975. Cognitive reference points. Cognitive Psychology 7.532–47.
ROSS, JOHN R. 1979. Where's English? Individual differences in language ability and language behavior, ed. by Charles J. Fillmore et al., 127–63. New York: Academic Press.
SPENCER, NANCY J. 1973. Differences between linguists and non-linguists in intuitions of grammaticality-acceptability. Journal of Psycholinguistic Research 2.83–98.
TRIPLETT, N. 1897. Dynamogenic factors in pacemaking and competition. American Journal of Psychology 9.507–33.
WAGENAAR, W. A., and P. PADMOS. 1971. Quantitative interpretation of stress in Kruskal's multidimensional scaling technique. British Journal of Mathematical and Statistical Psychology 24.101–10.
ZAJONC, ROBERT B. 1965. Social facilitation. Science 149.269–74.

[Received 20 June 1979;
Revision received 14 May 1980;
Accepted 4 September 1980.]