

Taking Performance Seriously*

Ivan A. Sag
Stanford University

1992

1 Competence and Performance

Thirty-five years ago, Noam Chomsky introduced the field of linguistics to new mathematical tools drawn largely from Recursive Function Theory. These were exciting tools that imparted mathematical precision, perhaps for the first time in the history of Linguistics, to the enterprise of grammar construction. The cornerstone of Chomsky's new theoretical edifice, was the *linguistic transformation*, an analytic tool which led to new insights about a vast array of empirical phenomena left unanalyzed by previous linguistic traditions.

As the theory of transformations developed, so did Chomsky's conception of linguistic theory. Transformational grammar was accorded the status of a theory of idealized linguistic knowledge – linguistic *competence*, to be distinguished from the more general study of real language use (including matters of on-line language processing), termed linguistic *performance*. The relation between these two notions, as Chomsky (1965:10) has emphasized, is that "... investigation of performance will proceed only so far as understanding of underlying competence permits."

For all their initial descriptive success, however, linguistic transformations have proven rather intransigent. Psychological models based on transformational grammar, e.g. the so-called 'derivational theory of complexity', were discarded as their predictions failed to square with the rapidly evolving developments of transformational theory (see Fodor, Bever and Garrett (1974) for an overview). And to this day no one, to my knowledge, has ever successfully formulated methods for computing with transformational grammars of the sort linguists have proposed, especially within the realm of parsing.¹

*For helpful discussions, crucial examples and/or occasional arguments, I would like to thank Alex Alsina, Ingrid Deiwiks, Eva Ejerhed, Janet Fodor, Mel Greenlee, Jerry Hobbs, Ron Kaplan, Martin Kay, Michele King, Mark Liberman, John Nerbonne, Fernando Pereira, Stanley Peters, Carl Pollard, Dave Rumelhart, Stuart Shieber, Hans Uzskoreit and Tom Wasow, most of whom surely disagree with something I say here. I would also like to express my gratitude to Professor Carlos Martin-Vide for his efforts on my behalf. This research was supported in part by a gift to Stanford University from the System Development Foundation.

¹There are modern computer systems that claim to be implementations of transformational grammar of one sort or another, yet all of these, to my knowledge, add a crucial device, the *surface covering grammar* – a set of phrase structure schemata that characterize directly the possible surface phrase structures of the language. Such systems are all nontransformational systems in disguise.

Yet the promissory note Chomsky issued in 1965 is surely long overdue. Theories of linguistic competence should be serving as a basis for testable models of linguistic performance. And it is high time that competence theories be the basis of work in computer language processing. And there are modern grammatical frameworks that compete with transformational grammar as models of linguistic competence. These frameworks are all based on the notion of satisfaction of constraints rather than the notion of transformation of structure. In this paper, I will suggest that this property makes them ideal candidates for performance-compatible theories of linguistic competence.

2 Constraint-Based Grammar

In the late 1970's, a number of radical alternatives to the transformational paradigm began to appear within the field of Linguistics. Reacting to the inelegance, psychological implausibility and empirical shortcomings of transformational analyses of the time, various linguists began to explore new alternatives. In this, they were aided by others – certain computational linguists who were trying to improve upon the ATN's of that era. Some of this work involved the move of thinking denotatively (model-theoretically) about linguistic information; other important developments involved reconsidering approaches such as categorial grammar and phrase structure grammar, previously dismissed as inadequate for natural languages by Chomsky and others in the early 1960's. The new strain of research in nontransformational syntax enjoyed an astonishing analytic success with respect to some of the most complex linguistic problems (e.g. the interaction of unbounded dependency constructions and constraints on coordinate structures) that have to this day remained unsolved in the transformational tradition. In an important paper, Pullum and Gazdar (1982) showed further that all extant arguments against such approaches (Context-Free Phrase Structure Grammars in particular) that had been advanced by Chomsky and others were based on mathematical errors and mistaken assumptions about crucial linguistic data. The stage was set for the new linguistics – one that could provide performance-compatible competence theories – one that had natural ties to other research within the emerging interdisciplinary field of Cognitive Science.

By the early 1980's, the theoretical landscape in Linguistics had changed considerably. Alternatives to the transformational orthodoxy now included Generalized Categorial Grammar, Generalized Phrase Structure Grammar (Gazdar et al. (1985)), Functional Unification Grammar (Kay (1979)) and Lexical-Functional Grammar (Kaplan and Bresnan (1982)). As the 1980's unfolded, these frameworks were refined in important ways, leading to the development of Head-Driven Phrase Structure Grammar (HPSG - Pollard and Sag (1987, to appear), the closely related Japanese Phrase Structure Grammar (JPSG - Gunji (1987)), and Categorial Unification Grammars (Uszkoreit (1986), Zeevat et al. (1987)), among others. These frameworks share two properties that are noteworthy from the outset. First, they are all based on the notion of constraint satisfaction. That is, the method of recursive function application that is central to classical transformational theory has given way to the method of recursive equation solving.² Second, the theories provide descriptions that simplify the grammatical information associated with a single sentence; and this information

²Pullum and Zwicky (1991) put this matter somewhat differently, contrasting the string rewriting algebras of transformational theories with the category structure algebras of constraint-based theories.

is associated directly with sentences that are produced, rather than with remote structures that must undergo potentially destructive and complex processes in the transformational derivation of such sentences.

3 The Significance of Constraint Satisfaction

In an important sense, the constraint-based alternatives to classical transformational grammar have made considerable analytical and theoretical progress.³ The analytical progress I refer to is in fact threefold. First, solutions were found for previously unsolved empirical problems (e.g. the interaction of unbounded dependencies and coordinate structures). Second, the new solutions that were posited embodied a tremendous simplification in the data structures (the wholesale elimination of transformational derivations) that the performance theory must reckon with. And third, new standards of precision were attained, e.g. rigorous formulations of rules, lexicons, and general theoretical principles.

The theoretical progress, on the other hand, is the performance-compatibility that follows from the shift away from systems that transform structures toward systems that satisfy constraints. Precisely because of their success in providing a computationally promising take on the difficult empirical problems of natural language description, these constraint-based theories – the first progeny of generative and computational linguistics – have been pressed into service by computational projects the world over, sparking a renaissance of enthusiasm for linguistics in computational circles where once there was only contempt.

One reason for this enthusiasm is that it is much easier to see how to build algorithms for successful implementations. In fact, there is a rich array of algorithms that the constraint-based architecture makes possible, allowing these grammars to be directly embedded within a theory of language processing.

This is, of course, the strongest and most interesting hypothesis available (the so-called strong competence hypothesis). And if we take this goal seriously, then there are a variety of properties of on-line language processing that might be taken to inform the design of grammar. For example, even the most superficial observation of real language use makes plain the fact that language processing is typically highly incremental: speakers are able to assign partial interpretations to partial utterances (and quite rapidly, in fact). Thus if a theory of grammar provides linguistic descriptions that demonstrably fail to be incrementally processable, then we should reject that theory as incompatible with the known facts of processing.

One striking example of this is English ‘echo questions’, as illustrated in the following kind of dialogue:

³I should add that the modern-day incarnation of TG - so called Government Binding Theory - retains only a vestige of the original transformational insight - a transformational rule called Move α , which most practitioners (Chomsky is a notable exception) regard as but a metaphor.

(1) A: Señora Maria Consuelo Bustamante y Bacigalupo is coming to dinner tomorrow night.

B:

ght

o did you say is coming to dinner tomorrow ni

h

w

Here, speaker A, in all likelihood suspends word-by-word processing of speaker B's utterance somewhere in the indicated region, once (s)he has recognized that the remainder of B's utterance is a repetition of A's own utterance. What examples like this show is that *partial* linguistic knowledge (e.g. the partial linguistic analysis of *who did you*, *who did you say* or *who did you say is*) is deployed on-line in an incremental fashion.

Unlike classic transformational theories, where syntactic knowledge is defined in terms of rules that map fully specified grammatical representations of sentences (phrase markers) into other such fully specified representations, constraint-based theories characterize all grammatical properties of sentences in terms of specific constraints that linguistic structures must satisfy. A subset of the constraints that hold true of an entire utterance is true of any substring of that utterance. Thus a constraint-based theory provides an incrementally available characterization of the initial substring of an utterance like the echo question in this example in terms of partial linguistic description.

Similarly, we know that language processing is highly integrative – information about the world, the context, the topic at hand, and so forth is skillfully woven together with linguistic information anytime that real language utterances are successfully decoded. For example, it is the encyclopedic fact that books don't fit on atoms – integrated mid-sentence – that allows the correct modification of the prepositional phrase *on the atom* to be determined well before word-by-word processing of a sentence like (2) is complete.⁴

(2) After finding the book on the atom, Kim decided that the library really wasn't as bad as people had been claiming.

Human processors employ grammatical information about a single sentence flexibly, at some points suspending word-by-word application of that information to consult information about prior discourse.

Without such nonlinguistic sources of constraint, the interpretation of even the most mundane of utterances can become highly indeterminate. So profound, in fact, is this indeterminacy (and the concomitant reliance of language on situational information) that the very fact that communication is possible using natural language acquires an air of considerable mystery. Although we lack at present any well-developed scientific theory of how linguistic and nonlinguistic information are brought together in the resolution of interpretation, it is nonetheless clear that we must demand of a linguistic theory that its grammars provide partial linguistic descriptions of a sort that can be flexibly integrated with nonlinguistic information in a model of processing. Constraint-based theories of language, in representing linguistic knowledge in terms of logic-based partial descriptions, allow for an integrated processing regime that consults information of both kinds on an 'as needed' basis.

⁴(2) is an adaptation of an example of Graeme Hirst's (see Hirst (1987)).

In addition to the incremental and integrative nature of human language processing, we may also observe that there is no one order in which information is consulted that can be fixed for all processing situations. In fact, it is possible to make an even stronger claim. In examples like (3), early accessing of morphological information allows the cardinality of the set of sheep under discussion to be determined incrementally, and well before the world knowledge necessary to select the ‘fenced enclosure’ (‘el corral’) sense of *pen*, rather than its ‘writing implement’ (‘el bolígrafo’) sense.⁵

(3) The sheep that was sleeping in the pen stood up.

In (4), on the other hand, the relevant world information – that sheep might fit inside a fenced enclosure, but not inside a writing implement – seems to be accessed well before the relevant morphological information constraining the cardinality of the set of sheep.

(4) The sheep in the pen had been sleeping and were about to wake up.

What contrasts like these suggest, is that the information accessed in on-line language processing is typically made available in an order dictated by the process in question. In comprehending these examples, for example, a hearer accesses morphological information earlier in (3) and later in (4) precisely because the order of access is tied fairly directly to the order of the words being processed. Assuming then that it is the particular language process that will in general dictate the order in which information is consulted, a grammar – if it is to play the role, as we assume, of information that fits directly into a model of processing – should be unbiased as to order. Grammars that are to fit into realistic models of processing should be completely order-independent.

Finally, we know that linguistic information, in the main, functions with like effect in many diverse kinds of processing activity, including comprehension, production, translation, playing language games, and the like. By ‘like effect’, I mean that the set of sentences potentially producible by a given speaker-hearer is quite similar to, in fact bears a natural relation (presumably proper inclusion) to, the set of sentences that that speaker-hearer can comprehend. This might well have been otherwise. The fact that there is so close and predictable a relation between the production activity and the comprehension activity of any given speaker of a natural language argues strongly against any theory where production grammars are independent from comprehension grammars, for instance. Rather, this simple observation suggests that the differences between, say, comprehension and production should be explained by a theory that posits different kinds of processing regimes based on a single linguistic description - a process-neutral grammar of the language that is consulted by the various processors that function in on-line linguistic activity.⁶

⁵I owe this sort of example to Martin Kay.

⁶The fact that production is more restricted than comprehension can then be explained within a theory of comprehension that allows certain kinds of linguistic constraints to be relaxed, or even word-by-word processing to be suspended when situational information is sufficient to signal partial communicative intent. Suspension of word-by-word processing clearly cannot enter into production in the same way (though incomplete sentences sometimes achieve communicative success). Hence, if we appeal to differences of process - not differences of grammar, there is at least the beginning of a natural account for why production should lag behind comprehension. Speakers that stray very far from the grammar of their language run serious risk of not being understood; yet hearers that allow grammatical principles to relax when necessary, will be understand more than those that don't. There is thus a deep functional motivation for why the two kinds of processing might differ as they appear to.

Observations of this sort about real language use and on-line language processing are quite robust, yet of course we could insist, as Chomsky has always insisted, that competence grammars are not to be directly embedded within models of performance (processing). With the new look of constraint-based grammars, however, it becomes possible to explore explicit models of processing where the linguists' grammars are directly embedded. As noted above, this is scientifically by far the more interesting hypothesis, and it is this very same possibility that gives new hope for the usefulness of constraint-based grammars for emerging language technologies.

In summary, grammars whose constructs are truly process-neutral hold the best hope for the development of processing models.⁷ And the best known way to ensure process-neutrality is to formulate a grammar as a declarative system of constraints. Such systems of constraints fit well into models of processing precisely because all the information they provide is on an equal footing. Comprehension models based on transformational grammar, by contrast, seem ineluctably saddled with the problem of *applying transformations in reverse*, and this is a problem that no one, to my knowledge, has ever solved.

The new kind of grammars, however – declarative in their formulation – exhibit no biases toward one mode of processing rather than another. Because each partial linguistic description is to be viewed denotatively, i.e. as being satisfied by a certain set of linguistic structures, the constructs of such grammars (e.g. words, rules, or principles) can be consulted in whatever order a process may dictate – the constructs are all constraints which allow themselves to be processed in a monotonic fashion. Given the current state of our knowledge of real-time language activity, a constraint-based architecture of this sort would seem to be the most plausible choice for the design of the theory of language, at least if the goal of embedding that theory within a model of language processing is ever to be realized.

One might try to cast a transformational grammar of the standard variety as a system of constraints, just as many procedural theories can be declaratively recast. The declarativization of transformational grammar would presumably involve taking transformational derivations (sequences of phrase markers) as the objects that linguistic descriptions describe. But the trouble with this attempt to make classical transformational grammars declarative is, apparently, the complexity of the data structures that it entails. In the absence of results about how to manage the complexity of this bloated space of objects (as Quine might have put it), it is difficult to see how this reinterpretation of the theory of transformations would bring us closer to understanding the facts of processing.

I am not saying that fully developed theories of processing already exist for constraint-based grammars (though such research has been ongoing for almost two decades). I am merely saying that grammars of this sort hold the greatest promise for such. They are highly constrained and have design properties (flexible, partial, order-independent, process-neutral descriptions that can be processed monotonically) which favor the development of explicit models of incremental, integrated processing. They are hence much more likely to play a role in explaining human language processing and to be useful in NLP.

⁷This point is made in passing already by Halvorsen (1983). See also Sag et al. (1986) and Fenstad et al. (1987).

4 The Resolution Problem

I have qualified my assessment of the success of constraint-based linguistics because I see an essential next step that must be taken, if we are to address the fundamental problem facing current research in natural language processing. This fundamental problem, which I will refer to as the *Resolution Problem*, arises from the simple fact that (as Barwise and Perry (1983) put it) ‘meaning underdetermines interpretation’. As is evident from the observations made in the previous section, the system of linguistic constraints that together constitute a natural language do not fully determine the interpretation of an utterance of any given sentence of that language. Linguistic knowledge of diverse kinds is smoothly and more or less effortlessly integrated with encyclopedic and situational knowledge by human language users in the process of communicating. The latest psycholinguistic studies, for example, are demonstrating massive integration of nonlinguistic knowledge and knowledge of the linguistic properties of an expression within a few hundred milliseconds after that expression is heard. Yet at present we have only the beginnings of a basic scientific theory of how this integration takes place.

The problem then is to provide a basic scientific answer to the following question:

(5) **The Resolution Problem for Natural Language Processing:**

How are diverse kinds of linguistic and non-linguistic knowledge integrated in on-line processing? (How is communication possible?)

To see the extent of this problem, which pervades all of language use, consider the following taxonomy of linguistic phenomena that give rise to what we might refer to as *communicative uncertainty*. First there are massive ambiguities of diverse sorts in all human languages:

(6) **Ambiguity:**

- Structural Ambiguity:

I forgot how good beer tastes. (beer in general, or good beer?) *I saw the man with the telescope.* (the man with the telescope, or saw with the telescope?)

- Lexical Ambiguity:

They can build a beter pen. (writing implement (‘bolígrafo’), or fenced enclosure (‘corral’)?)

The robot wouldn’t run. (wouldn’t move quickly (‘no corria’), or wouldn’t function at all (‘no funcionaba’)?)

- Ambiguity of Scope:

Jones found a defect in every car with over 500 miles. (one recurring defect, or a different one for each car?)

Everyone in the room speaks at least two languages. (the same two languages, or possibly different ones?)

- Ambiguity of Ellipsis:

Jones likes Smith more than Parker. (more than Parker does, or more than Jones likes Parker?)

Linguists have succeeded in providing detailed taxonomies of such ambiguities. Furthermore, now there are beginning to be constraint-based grammars that provide detailed accounts of the intricacies of these phenomena.

Second, there are many linguistic phenomena that contain *contextual parameters* – essentially pointers to aspects of the context that supply information essential to understanding the interpretation of the utterance:

(7) **Uncertainty of Reference:**

He is crazy. (Who is he?)

John is in charge. (John who? in charge of what?)

She ran home afterwards. (Who is she? whose home? after what?)

The relation is unclear. (relation between what and what? unclear to whom?)

(8) **Uncertainty of Relation:**

The nail is in the bowl. (nailed into the bowl, or resting inside of it?)

John's book (the book John owns?/wrote?/edited?)

The Amsterdam book (the book about Amsterdam? in Amsterdam? that was first discovered/read in Amsterdam?)

Researchers in linguistic semantics have also developed reasonably successful constraint-based frameworks (e.g. Situation Semantics) for the analysis of such 'parametric' aspects of interpretation.

Third, there are many linguistic phenomena, some bordering on what is traditionally classified as 'metaphor', where the literal interpretation of an expression is made vivid in diverse ways, depending on context, or else 'coerced' into a related interpretation:

(9) **Vivification (general meanings narrowed in context):**

Craig cut the lawn/hair/cocaine/record/rookie. (what kind of cutting?)

Coffee? (The rising intonation conveys either:

'I am tentatively suggesting an answer to your question', e.g. 'What used to be Columbia's most valuable cash crop?')

'I'm asking whether you want some coffee', or 'I'm asking you whether this is coffee.')

(10) **Coercion:**

The Boston Office called. (The intended interpretation is that someone associated with the Boston office called.)

These problems have been studied in depth by researchers in various linguistics traditions, including computational linguistics. in AI. Examples of this kind have been modelled in a successful, yet preliminary way by the work on *abductive reasoning* conducted by Hobbs, Stickel, Charniak, Goldman and others. (See, for example, Hobbs et al. (1990).)

Finally, various cognitive scientists have also studied and developed models of a fourth type of communicative uncertainty, which we may refer to as 'uncertainty of import':

(11) **Uncertainty of Import:**

I thought Jones was a spy. ('I was right all along.' vs. 'I was mistaken.')

Smith has outstanding penmanship. (praise in a letter of recommendation for a calligrapher; the kiss of death in a recommendation letter for a philosopher)

This phenomenon, which can be thought of simply as 'reading between the lines', involves drawing inferences based on recognition of the speaker's plans and goals. Preliminary computational models of just such reasoning have been developed by a number of researchers working on the boundaries of linguistics, AI and Cognitive Psychology.

In short, we now have a reasonably clear picture of the ways in which language gives rise to communicative uncertainty – a picture of the space in which resolution takes place. We also know something about the kinds of reasoning that must be performed in order to model successful communication. The important challenge facing research in natural language processing is to bring together the various kinds of research that have achieved this understanding in such a way as to take the next important step, to develop a basic scientific theory of the Resolution Problem.

Now one might ask whether we really need to have such a theory in order to get on with the business of developing computer language processing technology. Translation, for example, seems in many cases to preserve ambiguity, as in examples like (12).

- (12) a. Our knowledge of the universe comes from studying electro- magnetic radiation emitted by heavenly bodies.
- b. Notre connaissance de l'Univers provient de l'étude des rayonnements électromagnétiques émis par les corps célestes.
- c. Nuestro conocimiento del universo proviene del estudio de la radiación electromagnética emitida por los cuerpos celestiales.

But this is not always the case. Consider the example in (13), for instance.

- (13) a. He looked at the woman with penetrating eyes.
- b. Il a regardé la femme de ses yeux pénétrants. Il a regardé la femme aux yeux pénétrants.
- c. Miro a la mujer con ojos penetrantes. Miro a la mujer de ojos penetrantes.

Here, if the prepositional phrase modifies the verb (and hence it is *his* eyes that are penetrating) then the preposition is *de* in French – *con* in Spanish; if the prepositional phrase modifies the noun, however, the preposition of choice will be *à* in French and *de* in Spanish. Hence a machine translation system that failed to assign the right parse to the input language sentence would run a serious risk of conveying incorrect information - of failing to perform its intended function.

Such cases are by no means isolated. Consider another example of English ambiguity:

- (14) a. I forgot [[how good] [beer tastes]].
- b. I forgot [how [[good beer] tastes]].

Here the different parses can affect translation in subtle, but unmistakable ways. In German, we see inflection only when the adjective *good* modifies *beer*. Thus (15a) translates only (14a) and (15b) translates only (14b).

- (15) a. Ich habe vergessen wie gut bier schmeckt.
- b. Ich habe vergessen wie gutes bier schmeckt.

And translating from English into French or Spanish, on the other hand, the two English parses produce entirely different locutions:

- (16) a. J'ai oublié comme la bière a bon goût. [= (15a)]
 - b. J'ai oublié quel goût a la bonne bière. [= (15b)]
- (17) a. Me olvidé lo bien que sabe la cerveza. [= (15a)]
 - b. Me olvidé de como sabe la cerveza buena. [= (15b)]

The point is clear – the Resolution Problem is critical for natural language processing. And the discrete methods of constraint-based linguistics, promising though they may be for the theory of performance, have only begun to make inroads into the problem.

5 The Future

How then, should research on language proceed? The key, I believe, is to bring together diverse avenues of inquiry in an effort to understand how logic, association, and activation are brought together in natural language resolution. Several basic, yet important observations about resolution phenomena suggest that certain existing approaches embody properties that will ultimately be valuable to us in our attempt to develop models of on-line resolution.

First of all, it seems to be the case that certain resolutions – certain word senses, certain choices of modifier attachment, perhaps certain coreference assignments, etc. – are stronger than others. Thus the 'bolígrafo' sense of the English word *pen* seems inherently stronger, more frequent, and so forth than the 'corral' sense of *pen*. Similarly (as recognized at least since Kimball (1973)), there is a statistical tendency for ambiguous prepositional phrases to modify the closest possible element, as in the examples in (18).

- (18) a. Pat figured that Lee wanted to take the cat out.
- b. Chris seemed nice to them.

In (18a), *out* is more likely to be construed with *want* than with *figure*, though both construals are in principle possible. The *to*-phrase in (18b) is similarly more likely to modify *nice* than *seem*. These statistical facts have also been claimed to mirror default-context preferences that can be demonstrated experimentally. Finally, pronouns tend to corefer with compatible intrasentential NP's, as in (19).

(19) Clarke_i said that was the music from his_i wedding

But we know that all such effects are only preferences that can give way to other factors when utterances are properly situated. To see that this is the case, consider the simple case of lexical disambiguation in (20) and (21):

(20) The entire store was in disarray. The pencils were unsharpened. The pens were empty.

(21) The entire ranch was in disarray. The barns were unpainted. The pens were empty.

In a context like (20), the ‘bolígrafo’ sense of *pen* is correctly selected within a fraction of a second, while in (21), it is the ‘corral’ sense of *pen* that is chosen as the appropriate ‘resolution’ of the utterance.

Intuitively, it is obvious why this should be so. In (20), there are strong associations between the meanings of the words in the prior context (e.g. *store* and *pencil* and the ‘bolígrafo’ sense of *pen*). Similar associations exist between the meanings of the words in (21) and the ‘corral’ sense of *pen*. These effects are just the kind of phenomenon that associationistic models do well, and these models should be explored, as they have begun to be, with an eye toward reconciling them with rule-based systems, possibly even trying to make the linguist’s rules emerge from the patterns of activity in an association-based regime. It is unfortunate indeed that tirades such as the one in Pinker and Prince (1988) have caused a barrier to be erected between the field of linguistics and the world of connectionist research.

In addition to association, there are many cases where it seems that some kind of deduction is playing a role in resolution. Consider the indicated noncoreference of an example like (22), which differs from (19) only in that *wedding* has been replaced by *funeral*:

(22) Clarke_i said that was the music from his_i funeral.

Here it seems that we reason rather quickly, making crucial use of the fact that *dead men don’t talk*, to the conclusion that *Clarke* and *his* do not corefer.

Similarly, consider our earlier example of structural ambiguity in the two distinct contexts indicated in (23).

(23) I just got back from Saudi Arabia.

I forgot how good beer tastes.

(24) I just got back from Pilsen.

I forgot how good beer tastes.

The Saudi Arabia context clearly favors the [[how good][beer...]] parse, while the Pilsen context creates a strong bias in favor of the [how [[good beer]...]] structure. To explain these observations, it seems that we must make reference to a reasoning process where the hearer tries to construe the first sentence of the two sentence discourse as a justification for the second. The structure is resolved in such a way as to provide, together with relevant knowledge of the world, a plausible way of establishing an intended justification. For example, the fact that Pilsen is famous for good

beer presumably activates the [how [[good beer] ...]] parse, whose interpretation fits into a chain of reasoning involving the speaker having forgotten the taste of good beer prior to going to Pilsen. On the other hand, the generally known fact that alcohol is prohibited in Saudi Arabia presumably activates the concept of alcohol. Because the issue of how beer (in general) tastes is a subcase of the issue of how alcohol in general tastes, this active concept presumably creates a bias toward the [[how good][beer tastes]] parse. This time, however, the chain of reasoning involves the speaker drinking beer after returning from Saudi Arabia.

And of course there is a similar feeling of deduction in such classic examples as (25), due originally to Karen Sparck-Jones.

- (25) a. The council refused the demonstrators_i a permit because they_i advocated violence.
b. The council_i refused the demonstrators a permit because they_i feared violence.

Association and activation seem to be working hand in hand with deduction here as well.

In developing models of linguistic competence based on satisfaction of constraints, rather than transformation of structure, we have taken an important first step toward developing a theory of linguistic performance, i.e. of making good on the promissory note issued by Chomsky a quarter century ago. To take the next step, and to develop along the way computer processing applications that go beyond the mere toys of today, we must refine our theories of competence. The new theories we develop should be such that they bring together rich linguistic description with the notions of inherent bias, inherent associations, ever-fluctuating degrees of activation, and logical deduction. We have waited too long for the theory of competence to inform the theory of performance; we must take performance seriously now as we develop the theory of competence.

Bibliography:

Barwise, Jon, and John Perry (1983). *Situations and Attitudes*. Cambridge, Mass.: MIT Press.

Chomsky, Noam (1965). *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.

Fenstad, Jens-Erik, Per-Kristian Halvorsen, Tore Langholm, and Johan van Benthem (1987). *Situations, Language and Logic*. Dordrecht: Reidel.

Fodor, Jerry, Thomas Bever and Merrill Garrett (1974). *The Psychology of Language*. New York: McGraw-Hill.

Gazdar, Gerald, Ewan Klein, Geoffrey K. Pullum, and Ivan A. Sag (1985). *Generalized Phrase Structure Grammar*. Cambridge, Mass.: Harvard University Press and Blackwell's.

Gunji, Takao (1987). *Japanese Phrase Structure Grammar*. Dordrecht: Reidel.

Halvorsen, Per-Kristian (1983). Semantics for Lexical-Functional Grammar. *Linguistic Inquiry* 14.4: 567-615.

Hirst, Graeme (1987). *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge: Cambridge University Press.

- Hobbs, Jerry, Mark Stickel, Doug Appelt, and Paul Martin (1990). Interpretation as Abduction. Technical Report 499. Menlo Park, California: SRI International.
- Kaplan, Ronald and Joan Bresnan (1982). Lexical-Functional Grammar: A Formal System for Grammatical Representation. In *The Mental Representation of Grammatical Relations*, edited by J. Bresnan. Cambridge, Mass.: MIT Press. pp. 173-281.
- Kay, Martin (1979). Functional Grammar. In Chiarello, C., ed., *Proceedings of the Fifth Annual Meeting of the Berkeley Linguistic Society*, pp 142-158.
- Kimball, John (1973). Seven Principles of Surface Structure Parsing in Natural Language. *Cognition* 2(1): 15-47.
- Pinker, Steven and Alan Prince (1988). On Language and Connectionism: analysis of a parallel distributed processing model of language acquisition. *Cognition* 28: 73-194.
- Pollard, Carl and Ivan A. Sag (1987). *Information-Based Syntax and Semantics, Volume 1: Fundamentals*. CSLI Lecture Note Series No. 13. Stanford University: Center for the Study of Language and Information. Distributed by University of Chicago Press.
- Pollard, Carl and Ivan A. Sag (to appear). *Agreement, Binding and Control: Information-Based Syntax and Semantics, Volume 2*. Stanford: Center for the Study of Language and Information and Chicago: University of Chicago Press.
- Pullum, Geoffrey K. and Gerald Gazdar (1982). Natural Languages and Context-Free Languages. *Linguistics and Philosophy*.
- Pullum, Geoffrey K. and Arnold M. Zwicky (1991). A Misconceived Approach to Morphology. *Proceedings of the Tenth West Coast Conference on Formal Linguistics*. Stanford University: Stanford Linguistics Association. In press.
- Sag, Ivan A., Ronald Kaplan, Lauri Karttunen, Martin Kay, Carl Pollard, Stuart Shieber and Annie Zaenen (1986). Unification and Grammatical Theory. In *Proceedings of the Fifth Annual Meeting of the West Coast Conference on Formal Linguistics*. Stanford University: Stanford Linguistics Association, pp. 238-254.
- Uszkoreit, Hans (1986). Categorical Unification Grammars. In *COLING-86*. Bonn. [distributed by the Association for Computational Linguistics]
- Zeevat, Henk, Ewan Klein, and Jo Calder (1987). Introduction to Unification Categorical Grammar. In *Edinburgh Working Papers in Cognitive Science, Volume 1: Categorical Grammar, Unification Grammar and Parsing*, edited by J. N. Haddock, Ewan Klein and Glyn Morrill. Edinburgh: Edinburgh University Press.