Unification and Grammatical Theory

(Article begins on next page)

UNIFICATION AND GRAMMATICAL THEORY

IVAN A. SAG (STANFORD UNIVERSITY)
RONALD KAPLAN (XEROX PARC)
LAURI KARTTUNEN (SRI INTERNATIONAL)
MARTIN KAY (XEROX PARC)
CARL POLLARD (STANFORD UNIVERSITY)
STUART SHIEBER (SRI INTERNATIONAL)
ANNIE ZAENEN (XEROX PARC)

1. Introduction

This paper informally presents a new view of grammar that has emerged from a number of distinct but related lines of investigation in theoretical and computational linguistics. Under this view, many current linguistic theories—including Lexical-Functional Grammar (LFG), Generalized Phrase Structure Grammar (GPSG), Head-Driven Phrase Structure Grammar (HPSG), and categorial grammar (CG)—fall within a general framework of UNIFICATION GRAMMAR. In such theories the linguistic objects under study are associated with linguistic information about the objects, which information is modeled by mathematical objects called FEATURE STRUCTURES. Linguistic phenomena are modeled by constraints of equality over the feature structures; the fundamental operation upon the feature structures, allowing solution of such systems of equations, is a simple merging of their information content called UNIFICATION.

Although differences among these theories remain great, this new appreciation of the common threads in research paradigms previously thought ideologically incompatible provides an opportunity for a uniting of efforts and results among these areas, as well as the ability to compare previously incommensurate claims.

Because of the brevity of the present work, we will necessarily be unable to present the formal underpinnings of unification grammar, relying instead on the reader's intuition to provide the necessary details. For more detailed discussions of the formal, computational and mathematical foundations of unification grammar, and its relation to some current linguistic theories, readers are referred to Shieber (1986) and works cited therein.

2. Derivational and Nonderivational Theories

The novel structure of unification-based descriptions of language is highlighted by contrasting it with earlier transformational descriptions. In the era of transformational grammar, linguistic theory was galvanized by the appreciation of the power and utility of manipulating structured expressions to provide a basis for linguistic description and explanation. By contrast, unification grammars describe language in terms of static constraints on information associated with structured expressions, as opposed to the dynamic transformation of the expressions themselves.

Within transformational linguistics, the syntactic structures relevant to human-language competence have been of two sorts: base structures specified in full by a phrase-structure component together with a lexicon, and derived structures produced by the application of transformational rules. In early incarnations of transformational grammar, where operations of movement, copying and deletion were freely employed, the various stages of derivation associated with one another through the application of transformational rules were quite diverse in nature. Phrases present at one level within a transformational derivation were at other levels either absent or in dislocated positions. Transformational derivations were opaque in the sense that information in corresponding pieces of structure at different stages of derivation are incompatible. As an example, we show in Figure 1 a standard deep structure for the noun phrase

(1) the person who Sandy was kissed by

After applying the passive transformation—involving deletion of the object NP, replacement of the subject, and insertion of a by-phrase—the intermediate structure shown in Figure 2 is derived. In Figure 3, we overlay the

FIGURE 1. Deep structure for Sentence 1

FIGURE 2. Intermediate stage after passive

FIGURE 3. Overlaying first two stages in derivation

two structures, showing that information in corresponding locations in the two structures is incompatible. For instance, in one structure, the lexical item *Sandy* occurs where *who* occurs in the other. Thus, in that position, such properties as definiteness of the noun phrase differ in the two stages. In other examples, incompatibilities of number, person, or other features
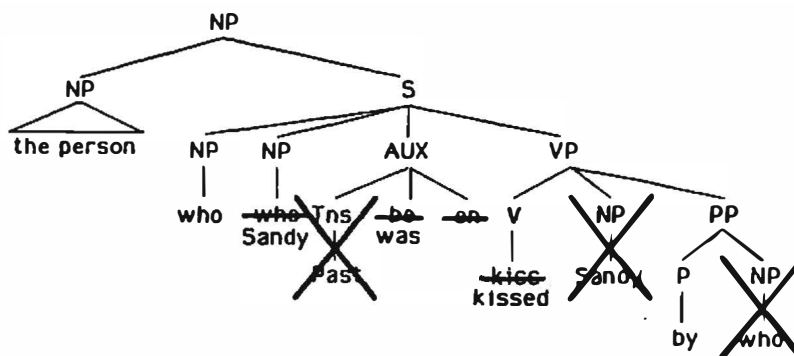
Fɪɢᴜʀᴇ 4. Overlaying all stages in derivation

can occur. Finally, in Figure 4 we show the overlaying of the stages in the
derivation of the final surface structure for the sentence. The applications
of the relativization and 'affix hopping' transformations further develop in-
compatibilities among the various stages of derivation.

By contrast, unification-based descriptions of language require applica-
tion of all linguistic constraints be ᴍᴏɴᴏᴛᴏɴɪᴄ, that is, the constraints merely
add information, without performing structural changes. If we were to
imagine a transformational theory with this property, its derivations would
be transparent; surface phrase markers would be essentially structurally iso-
morphic to all the phrase markers in their syntactic derivation.[1] For instance,
a derivation for (1) in such a theory might develop as in Figures 5 through 7.
Figure 5 shows a proposed deep structure for the sentence. The passive
constraint merely adds the information that two NPs are coindexed and the
lower one is phonetically unrealized. Adding this information yields the
structure in Figure 6. Adding the relativization constraint leads to the final
structure in Figure 7.

Note that in these stages of derivation, no incompatibilities arise in the
combination of one structure with another. Indeed, the entire derivational
process can be viewed as a process of adding compatible information to the
original structure, information such as indices, cases, and so forth. Also,

---

[1]Subsequent developments within transformational theory, such as the structure preserving
hypothesis, trace theory, strong versions of the lexicalist hypothesis and the projection
principle, suggest that transformational grammar may be moving in this direction.

FIGURE 5. Alternate deep structure for Sentence (1)

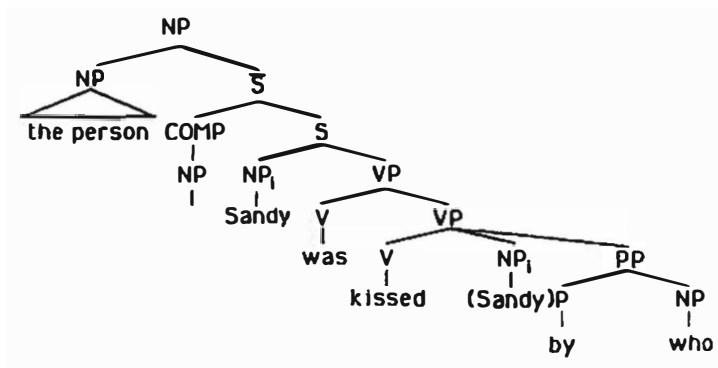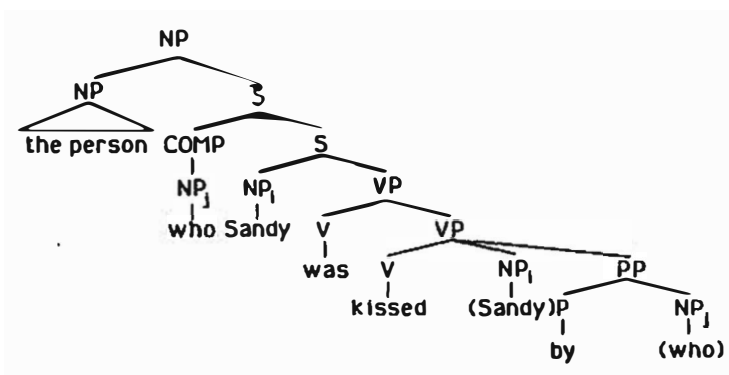FIGURE 6. Overlaying intermediate stage after passive

FIGURE 7. Overlaying all stages in derivation

the order of application of the constraints, in this example and in general, is purely arbitrary; other orders yield the same results. Therefore, in such

a theory, derivation-final structures by themselves can provide the information required for determining well-formedness and the system has no need for syntactic derivations. In such a system, syntactic rules, lexical entries, universal principles and language-particular parameters can all be viewed as simultaneous constraints on output structures—constraints on indexing, on 'binding', or on whatever kinds of grammatical information turn out to be appropriate to associate with structured expressions.

In addition, the relevant constraints used in linguistic theories involve *equality* of information, such as the identity of indices associated with distinct NP's in specifiable configurations or domains. This is especially important, for, as we shall see, systems of equality constraints are monotonic in the sense just described. In summary, unification-based theories, in contrast to early transformational systems, embody a nonderivational view of linguistic theory which countenances structured expressions (represented by parse trees) and a set of constraints (some universal and some language-particular) which impose conditions of equality on the grammatical information associated with various parts of those expressions.

3. Unification Grammar through Examples

We now turn to a fuller (though still abridged and informal) discussion of the unification grammar framework, presenting in more detail key concepts of *equality* of and *partiality* of information and of *unification* through a series of examples.

3.1. **Equality and Partiality of Information.** Let us begin with a simple example of subject-verb agreement. In English, both finite verb forms and noun forms may bear information about the person and number features of the subject of a sentence of subject-predicate form. These two pieces of information are subject to a condition of EQUALITY that is entailed by any descriptively adequate theory of agreement. This condition is what guarantees that the sentences in (2), but not those in (3), are grammatically well-formed.

(2) a. The building is shaking.
     b. The buildings are shaking.

(3) a. * The building are shaking.

    b. * The buildings is shaking.

But it is an inevitable fact of human language that linguistic expressions bear partial information. *Partiality* of information is used in many linguistic analyses to eliminate the need for specifying a set of alternatives repeatedly. For instance, rather than specifying in numerous places that adjective and noun phrases behave in similar ways, we might decompose the category symbols into separate features for N and V, as in versions of X-bar theory, and pick out the class containing APs and NPs by a partial information structure that specifies a + value for N but no value for V. Such decompositions to eliminate the proliferation of fully-specified structures are commonplace in linguistic analyses.

In the case of agreement being discussed here, partiality is evident in that the verb phrase may contain no information whatsoever about the subject's agreement features, as in (4).

(4) The building had been shaking.

Or the subject NP itself may bear no such information, as in an example like (5):

(5) The salmon have been jumping.

In this case, a constraint requiring equality between agreement information on subject and verb would still be solvable under the assumption that the subject is plural. The same solution method, then, can be used to infer the number of the subject in (5) and the verb in (4). Again we contrast this with earlier transformational systems, in which these inferences would derive from homophony of two forms of *had* and *salmon* differing only in their number. Besides introducing artifactual directionality and ordering into a system which, as we shall see, does not require it, such a system requires postulation of extraneous lexical entries for nouns, verbs, and other lexical items. The profligacy of homophonous forms is even more prevalent in languages with richer morphology. In any case, the use of partiality of information (as in featural decomposition) is traditionally recognized as an appropriate technique for capturing this type of classificatory generalization.

The observation enabling the unification grammar view is that the two concepts just presented—equality and partiality of information—are not only obviously necessary for inclusion in a linguistic system, but by and large sufficient for capturing syntactic phenomena. One requirement to substantiate this claim is the ability to combine constraints to model more complex phenomena. For instance, in more complex cases in which the agreement constraint just described interacts with relativization, such as

(6) a. The salmon which has been in the lake has been jumping.
    b. The salmon which have been in the lake have been jumping.
    c. * The salmon which have been in the lake has been jumping.
    d. * The salmon which has been in the lake have been jumping.
    e. The salmon which had been in the lake has been jumping.
    f. The salmon which has been in the lake had been jumping.

it remains to be demonstrated that the same agreement constraint suffices to predict grammaticality.

Viewed schematically, these sentences all manifest a structure in which the agreement information on the subject's head noun and on the form of *have* in the matrix are constrained to be equal (as in the previous examples). Furthermore, the head noun and the *have* form in the embedded clause are constrained to be equal as well, either by virtue of an intermediate trace in the embedded clause (as in traditional LFG analyses) or directly (as in GPSG). Finally, the various verbs and nouns may individually contribute full or partial agreement information as a further lexical constraint. Solving these constraints will yield the grammaticality distribution in (6), for even though the subject head noun is unmarked for number, it is a consequence of the transitivity of equality that the agreement information associated with the two forms of *have* will be equal. In (6c) and (6d) this constraint is violated. This conclusion holds for any unification grammar, since all describe these constraints in terms of static equalities. The particular source of the constraints differs from one theory to another, and this constitutes an important theoretical distinction among the theories. The order-independence

of equation-solving, however, guarantees that the source of the equations—whether lexical or syntactic, universal or language-particular—will have no bearing on their solutions.

Thus, in this simple case at least, the equality constraints for different syntactic phenomena (i.e., subject-verb agreement and relativization) interact properly in a way that is order-independent. In Section 4, more complex examples of interaction will be considered.

3.2. **Unification.** As we have seen, stating constraints as equality conditions over partial information structures is a powerful method of describing linguistic phenomena. Because solution of such systems of equations is the primary tool in unification grammars, a crucial criterion of the adequacy of such systems is the existence of solution techniques for such equations. Indeed, the primary technique for solution of equality constraints lends its name to the entire paradigm, for it is the operation of UNIFICATION itself.

Unification is an operation that does nothing more than to amalgamate compatible partial information and to fail to amalgamate incompatible partial information. For example, suppose we encode the agreement information associated with a third-person noun phrase (like *salmon*) with the feature structure

$$(7) \quad \begin{bmatrix} cat: & \begin{bmatrix} n: & + \\ v: & - \end{bmatrix} \\ person: & third \end{bmatrix}$$

The constraint that the subject of the finite verb *have* is a plural NP would be conveyed in terms of the feature structure

$$(8) \quad \begin{bmatrix} cat: & \begin{bmatrix} n: & + \\ v: & - \end{bmatrix} \\ number: & plural \end{bmatrix}$$

Combining these feature structures by unification, we would arrive at the feature structure encoding the sum of the information content of these two feature structures, namely,

$$(9) \quad \begin{bmatrix} cat: & \begin{bmatrix} n: & + \\ v: & - \end{bmatrix} \\ number: & plural \\ person: & third \end{bmatrix}$$

Of course, not all feature structures can be combined by unification. Consider the subject requirements of the verb *has* which we might summarize as

$$(10) \quad \begin{bmatrix} cat: & \begin{bmatrix} n: & + \\ v: & - \end{bmatrix} \\ number: & singular \\ person: & third \end{bmatrix}$$

An attempt to unify (10) with (8) (as would be required in any unification-based analysis of (6c) or (6d)) could not succeed, as there is no feature structure which contains both the information that the number is singular and that the number is plural. Unification is said to FAIL in this case, accounting for the ungrammaticality of (6c) and (6d).

In just this way, when we want to identify two linguistic constructs about which we have only partial information encoded in feature structures, we use equality statements in formulating the appropriate linguistic principle or constraint, and these equalities can then be solved by unifying the corresponding structures.

4. ANALYSES USING UNIFICATION

In this section we present two examples of unification-based analyses of linguistic phenomena, exemplifying two broad classes of linguistic phenomena: long-distance syntactic dependencies and lexical dependencies. We show that the examples can be handled using the same techniques as were introduced in the previous section. The informal analyses presented here are therefore independent of any particular linguistic theory, but highly dependent on a unification grammar setting. The differing nature of the classes of phenomena to which these examples belong should indicate the

FIGURE 8. Schematic representation of dependencies in Germanic filler-gap constructions

scope of application allowed by unification grammar techniques in encoding linguistic generalizations.

4.1. **Germanic Unbounded Dependencies.** The first analysis concerns unbounded filler-gap dependencies in Germanic languages which exhibit various grammatical restrictions (e.g., choice of grammatical category, case or other inflectional parameter) holding between a filler and, for example, a verb governing the trace bound by the filler. There is essential agreement across grammatical frameworks that such restrictions are to be analyzed in terms of two simultaneous dependencies: one holding between the verb and a phonologically unexpressed object (or trace), and another holding between the filler and its trace, as illustrated in Figure 8.

Because the syntactic information the verb requires of its object is identified with that of the phonetically unexpressed element, which in turn is identified in relevant respects with the filler that binds it, the information borne by the filler must be compatible with the verb's requirements. Unification of three distinct pieces of syntactic information is the essence of such analyses, however formalized, including those presented in Kaplan and Bresnan (1982), Gazdar et al. (1985), and Pollard (in press). Thus, any analysis within this paradigm is guaranteed-by virtue of the monotonicity of equality systems-to scale up to more complex interactions with other constructions.

4.2. **Modern Irish Verbal Forms.** A second example involves a lexical dependency turning on differences between analytic and synthetic verbal forms in Modern Irish as discussed by McCloskey and Hale (1984) and Andrews (1984). Analytic verb forms in Irish (such as *chuir*, the past tense form of the verb meaning 'to put') require subjects whose lexical head is overtly expressed, whereas synthetic verbal forms (e.g., the 1st person singular conditional form of the same verb, *chuirfinn*) require subjects whose head is unexpressed, as illustrated in (11).

(11) a. Chuir   me fein isteach ar an phost sin
         put(PAST) I   EMPH in   on the job   that

      b. Chuirfinn fein isteach ar an phost sin
         would put EMPH in   on the job   that

      c. * Chuir   fein isteach ar an phost sin
           put(PAST) EMPH in   on the job   that

      d. * chuirfinn   me fein isteach ar an phost sin
           would put I   EMPH in   on the job   that

Unification-based analyses of this set of data (and many others like them) have been developed within LFG and HPSG. In LFG, the facts are dealt with in terms of f-structure consistency (that is, solvability of the system of equations) and certain independently motivated assumptions about completeness of f-structures. Figure 9 illustrates the constituent structure for (11a) with associated f-structure projections. Two-headed arrows mark the equational constraints that would be imposed in the LFG analysis.

Because the lexical form *chuir* assigns no index to its f-structure subject, such an index must be supplied by a lexically expressed NP. Otherwise the resulting f-structure will be incomplete, in violation of universal principles. And as shown in Figure 10, a synthetic lexical form like *chuirfinn* does assign an index to its f-structure subject, hence any combination with an overtly expressed subject phrase projecting its own f-structure index (universally, all overtly expressed lexical forms bear distinct indices in LFG) results in unification failure—the equations are unsolvable. In the figure, incompatible feature values leading to the failure are highlighted by surrounding circles.
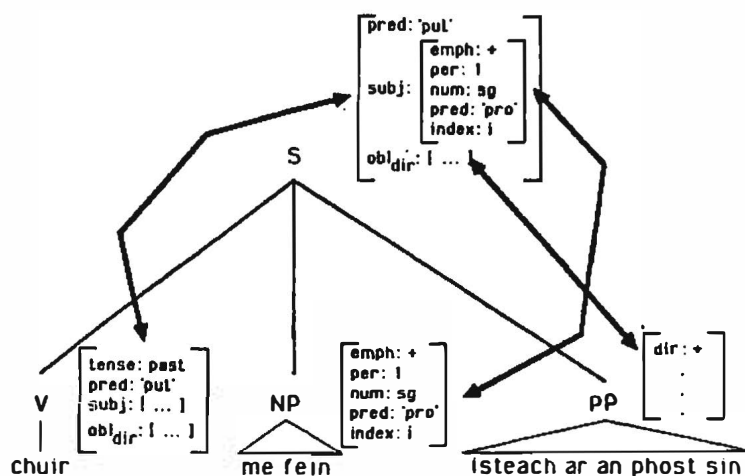
FIGURE 9. LFG account of (11a)

FIGURE 10. LFG account of (11d)

In HPSG, where subcategorization by a verbal head is treated by means of the list-valued feature *subcat*, members of the list must unify with the appropriate dependent elements (this is ensured by grammar rules and universal grammatical principles). The analytic verbal forms subcategorize for non-null subjects, NP's specified as [*nform*: *norm*], and hence combine with these (but not non-null subjects), as illustrated in Figure 11.
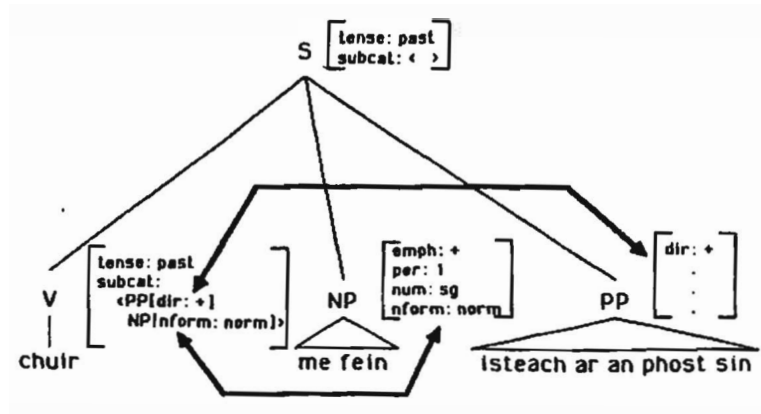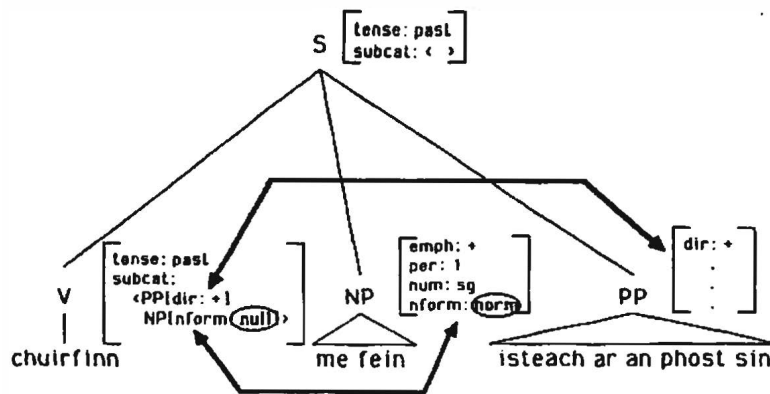
FIGURE 11. HPSG account of (11a)

FIGURE 12. HPSG account of (11d)

But synthetic verbal forms select headless subject NP's (those marked [*nform*: *null*]) in HPSG. Hence any combination of these forms with full subject phrases (all specified as [*nform*: *norm*]) results in unification failure, as illustrated in Figure 12.

Although superficially different, these two analyses have an underlying similarity in their reliance on the existence or nonexistence of solutions to systems of equations. Consider the ungrammaticality of (11d). In both analyses, the equations require identity of information about the subject as specified on the verb and subject. And both theories require different values for some aspect of the information on the verb and subject (index information in LFG and nform information in HPSG).

Now the critical equality follows from transitivity applied to two constraints in an LFG grammar rule, and from a principle of subcategorization in HPSG. The index and nform information are provided through a variety of mechanisms in the two systems. Yet the fact that both systems can be viewed as merely stating sets of equality constraints, and stating the same equality constraints in this case, together with the monotonicity of equality systems, suffices to show that both analyses will predict the same grammaticalities.

These analyses embody different hypotheses about how informational constraints should be decomposed into natural classes and how constraints interact with one another. And these are important theoretical issues whose resolution is the object of ongoing research. Yet all of the frameworks just illustrated exhibit profound substantive similarities with respect to fundamental mechanisms. All embody the hypothesis that linguistically significant generalizations are to be expressed in terms of identity constraints superimposed on structured expressions rather than in terms of the derivational history of those expressions.

5. Why Unification Grammar?

Unification grammar methods have been used in the analysis of numerous varied linguistic phenomena. Several advantages over derivational methods have been alluded to in the foregoing discussion. In particular, stating linguistic constraints directly in terms of systems of equations frees us from making decisions as to 'direction of movement'. Having to make such decisions can in turn cause artifactual idiosyncrasies as in whether to copy agreement features from subject to verb or verb to subject. As another example, in describing extraction from conjoined phrases, we must describe which of the two traces is the source of the filler and which is deleted, or alternatively, invent a method for combining traces (checking for compatibility of information) before moving the combined structure to the filler position. Of course, this latter alternative is quite likely a reinvention of unification. The former method is heir to well-known difficulties.

Another artifact of derivational analyses is the fact that phenomena such as the subject-verb agreement in examples (4) and (5) require a proliferation of homophonous lexical entries.[2] The natural incorporation of partial information into unification-based systems frees us from postulating multiple fully-specified homophonous lexical items in these cases.

For computational reasons, monotonic systems have some advantages over derivational ones. For instance, in attempting to implement parsers for early transformational systems, the order-dependency of transformations required efforts to 'reverse' the transformations which proved a difficult, if not hopeless task.[3] The existence of unification as a simple, order-independent, computationally precise method for solving any system of equations of the sort used in unification grammars allows us to build very general interpreters for unification grammars that can be used to test analyses in many of the different unification-grammar theories. There is no need to reverse operations.

A theory of grammar that allows many different implementations, as monotonic theories do, enlarges the domain in which psychological processes are free to operate, allowing a direct embedding of the theory of linguistic knowledge within a reasonable model of language processing. There is every reason to believe that diverse kinds of language processing—syntactic, lexical, semantic and phonological—are interleaved in language use, each making use of partial information of the relevant sort. Given that this is so, the theories of each domain of linguistic knowledge should be nothing more than a system of constraints about the relevant kind of linguistic information—constraints that are accessed by the potentially quite distinct mechanisms that are involved in the production and comprehension of language.

The well-understood semantics of the formalisms we employ enables us to achieve a precision in linguistic science that has been conspicuously absent in many recent debates. For the first time in recent memory, it becomes

---

[2]Moving the homophony problem into a set of 'spell-out rules', of course, only postpones the problem. The statement of the rules themselves would be prone to the same proliferation unless partiality of information were used.

[3]The results of Peters and Ritchie (1973), in fact, show that the job is computationally undecidable.

possible to synthesize in a rigorous way results obtained within divergent research traditions. The development of this conceptual framework in a mathematically precise manner has enabled us to systematically compare proposals in a number of seemingly diverse linguistic frameworks, finding communality as well as clarifying important differences, in short distinguishing between matters of *notion* and matters of *notation*.

6. CONCLUSION

In the course of the last few years, we have come to the realization that much of current linguistic practice in many of the rival theories can be viewed from a single unifying perspective, and have been led to search for a general conceptual framework in which to cast proposals made in any number of differing linguistic theories, such as LFG, GPSG, HPSG, CG, as well as closely related work on computational linguistics in such frameworks as FUG (developed at Xerox PARC) and PATR (developed at SRI International). Work done in these theories has made various assumptions about the nature of the information manipulated by grammatical constraints as well as differing assumptions about the nature of the constraints themselves. But the development of a common general framework, unification grammar, has enabled us to isolate theoretical communality, integrate analytic techniques and to clarify the nature of theoretical controversies. It seems to be an emerging consensus of modern linguistics that explanatory accounts of syntactic phenomena can be provided in a monotonic system of equality constraints over partial information structures associated with structured expressions.

REFERENCES

Andrews, Avery. 1984. Lexical insertion and the elsewhere principle in LFG. Manuscript.

Gazdar, Gerald, Ewan Klein, Geoffrey K. Pullum, and Ivan Sag. 1985. Generalized phrase structure grammar. Oxford: Blackwell.

Kaplan, Ronald and Joan Bresnan. 1982. Lexical-functional grammar: a formal system for grammatical representation. The mental representation

.. of grammatical relations, ed. by Joan Bresnan, 173-281. Cambridge: MIT Press.

McCloskey, James and Kenneth Hale. 1984. On the syntax of person-number inflection in modern Irish. NLLT 1.487-534.

Peters, Stanley, and Robert Ritchie. 1973. On the generative power of transformational grammars. Inf. Sc. 6.49-83.

Pollard, Carl. In press. A semantic approach to extraction in a monostratal theory. To appear in L&P.

Shieber, Stuart. 1986. An introduction to unification-based approaches to grammar. Stanford: Center for the Study of Language and Information.