Encylopedia of Cognitive Science

Head-driven Phrase Structure Grammar

Robert D. Levine
Associate Professor, Department of Linguistics

Ohio State University
Columbus, Ohio USA 43210–1298
levine@ling.ohio-state.edu

January 7, 2003

Head-driven Phrase Structure Grammar

1 Basic concepts

Generative theories of linguistic structure can be contrasted along various dimensions, the most fundamental of which divides them into theories which license sentences by constructing a set of structures and then deriving further structures from these, until some point where no further structures are derivable, versus theories which admit only a single structural object corresponding to each sentence (or a set of such objects, corresponding to syntactic ambiguity), where admission entails simultaneous satisfaction of a set of applicable constraints. The first approach, generally referred to as *transformational* grammar, typically uses the successive stages in the derivation of the sentence to establish appropriate linkages in information content among parts of the representation, while the second, often called *monostratal*, establishes such linkage by directly constraining information distribution, typically encoded as the respective values assigned to certain attributes, or, as I refer to below, *grammatical features*.

A second significant distinction grammatical divides frameworks into those whose admissible objects are characterized by *phrase structure*, versus those in which no persistent configurational properties of representations are posited. *Categorial* grammars, for example, specify functor/argument possibilities for individual lexical items, where such possibilities determine whether or not some sequence of words can combine to give rise to an object corresponding to the type associated with a clause, but the order of combination does not impose structural configurations on the the string or its substrings which are subsequently available for stating rules or generalizations.

A third important division distinguishes theories which require a one-to-one match between syntactic form and propositional structure from those which do not. In a sentence such as *Robin seemed to understand Dana's arguments*, for example, it is evident that the individual named Robin has a semantic relationship to the predicate denoted by the linguistic form *understand Dana's argument*, while having no semantic relation to the verb *seemed*. That this is generally true for the subject of such verbs (standardly referred to as *raising* verbs) is made clear by a variety of classical syntactic diagnostics (preservation of idiomaticity, preservation of truth under passivization, and so on). Yet in the syntax of this sentence, *Robin* is the subject of *seemed*, with no evident local structural relationship to *understand Dana's arguments*. There is thus a mismatch between syntactic form and semantic interpretation in such examples. On one standard approach to this problem, the mismatch is the result of a derivational history in which a structure with *Robin* appearing in the subject position of *understand Dana's arguments* in some early structural representation is displaced to the subject position of *seemed*. On the other approach, *Robin* is syntactically always and only a subject of *seemed*, but is linked through various constraints to the semantic content of *understand Dana's arguments*, without any comparable linkage to the semantics of *seemed*. The former approach will naturally be restricted to derivational theories, but not all derivational theories follow this approach; see, e,g., Brame 1976, Culicover and Wilkins 1984. On the other hand, nonderivational theories have no choice; they must take as basic the syntax/semantic mismatch exemplified here and motivate the discrepancy in their own formal structure—via meaning postulates, lexical entailments, representational properties, or some other means.

Against this background, phrase structure grammars can be concisely characterized as a family of theories which

- assume only a single level of representation;
- take this level to consist in a set of of hierarchically structured objects, in which
- syntactic and semantic properties need not mirror each other in any respect.

Within the family of PSG theories, there is a considerable variety of possibilities. In the following discussion I illustrate the possibilities of the overall PSG framework with reference to one particular version, Head-driven Phrase Structure Grammar (HPSG).

2 HPSG

2.1 Ontological assumptions

HPSG, developed by Carl Pollard and Ivan Sag over the past fifteen years (Pollard and Sag 1987, Pollard and Sag 1994) initially as a refinement and extension of Generalized Phrase Structure Grammar (Gazdar 1981, Gazdar et al. 1985), belongs to a family of phrase-structure-theoretic approaches in which a rich set of lexical specifications, coupled with a few very general combinatorial constraints and restrictions on information-sharing, interact monotonically to give rise to sets of complex objects called *feature structures*, which model the properties of linguistic signs (though certain recent developments of HPSG have abandoned strict monotonicy by allowing defeasible default characterizations of feature-sharing principle to interact with 'hard' lexical entries or type declarations; see discussion in section 3 below). Each feature structure admitted by an HPS grammar for some language licenses an expression in that language. Considerable effort has been made to provide an explicit model theory for HPSG grammars, and to formulate the constraint systems which HPS grammars themselves consist of as statements in certain formal logics. Such logics are comparable to familiar systems such as the first-order predicate calculus, with the crucial difference that constraint (non)satisfaction takes the place of model-theoretic (non)satisfaction of, for example, a set of conventional formulæ by some interpretation given for the latter. This approach, which has its origin in the computational linguistics of the mid-1980s, greatly expedites implementation of HPS grammars, but from the theoretical perspective its main significance is that the denotation of particular analyses and claims stated in the framework is always mathematically explicit, a necessary but often neglected condition for generativity.

An HPS grammar is a *description*, i.e, a set of formulæ in some feature logic which an arbitrary feature structure of feature structures either does or does not satisfy (see, e.g., Richter et al. 1999 for an example of such a logic that is becoming widely employed in HPSG), where only feature structures that satisfy the formulae of the grammer are taken to be complete representationd of the grammatical properties of some linguistic expression. Feature structures themselves are sorted unary partial subalgebras in which feature names are functions from the set of nodes to the set of nodes, with each node assigned a *type* or *sort* label. Each node corresponds to some object comprising certain properties, and the node which a given feature maps this object to corresponds to the value of that object for the property named by the feature.

Consider, for example, an auxiliary verb in English. Such verbs select a certain class of subjects and appear either in canonically ordered sentences (*Robin has left*) or inverted sentences (*Has Robin left?*). In HPSG, such a linguistic expression will be modeled by a feature structure including a specification for the feature CAT, providing relevant syntactic information, including a specification of the HEAD properties of that expression—those which are invariably shared between mother and head daughter. The feature HEAD is thus taken to be a function which maps a particular node labelled by the sort *category* to a particular node of sort *noun*, *verb* and so on; for verbs, this latter node itself is mapped by a function VFORM to a node labelled by one of of a set of sorts *fin(ite), inf(initive), b(a)se, ...*, by a function AUX(ILIARY to one of the sorts *plus* or *minus*, and so on. In the same way, the feature SUBJ is a function from nodes of sort *cat* to node whose type corresponds to a (singleton) list identifying the kind of subject the verb can combine with, and the feature COMPS a function from *cat* to a different node of the same sort, identifying the morphosyntactic properties of the complements that the verb selects; these features are discussed in greater detail below.

More generally, various versions of HPSG model theory essentially converge on the characterization of the model universe for an HPS grammar along the following lines:

- Each grammar specifies a *signature*—a set of feature names, a set of sort names, and an assignment function which pairs with each feature name a mapping from the set of sort names to the power set of sort names. The signature in effect defines what kinds of objects there are, and what kinds can be mapped to (i.e., specified for) what other kinds of objects under the operation denoted by some feature name.
- An *interpretation* for an HPSG grammar is a set of class of partial algebras, each of which specifies a set of sorted nodes and a set of unary operations, corresponding to the feature names, each of which takes some subset of nodes as its domain and some subset of nodes as its range, where for all mappings from subset to subset, the sort

labels of the domain and range elements must adhere to the constraints imposed by the assignment function.

- Given the objects and functions so far defined, it is possible to construct a formal notion of pathway—sequences of nodes—in terms of successive applications of appropriate feature-functions. Thus, the pathway CAT|HEAD|VFORM specifies a function which maps an input node of a certain sort to a node of a sort such as *fin*. Feature logics can be defined on such pathways, and an HPSG grammar specifies a number of *constraints*, stated in terms of such pathways. Constraints are descriptions that are satisfied for just those sets of partial algebras in which all pathways specify the values that the constraints of the grammar require them to. The constraints, as noted earlier, in effect play the role of meaning postulates in classical model theories of natural language semantics: they restrict the set of admissible interpretations of the grammar—in this case, sets of partial algebras—to just those which reflect the property of the natural language modeled by the grammar.

These components of the HPSG model theory express standard HPSG assumptions that the object which is the value of any given feature belongs to a particular sort, that only objects of certain sorts can be the value of any given feature, and that objects of different sorts are characterized by different sets of features respectively. The interaction of these assumptions constitutes an overall constraint on the possible set of feature structures that the grammar admits, corresponding to just the set of linguistic expressions taken to be well-formed by speakers.

It should be noted that this algebraic characterization of feature structures is similar to the formalization of category specifications of GPSG as presented in, e.g., Gazdar et al. 1985 and Gazdar et al. 1988. Discussions of feature structures, however, typically represent them not as partial algebras, but graphically—specifically, as pointed directed graphs in which nodes, each annotated with a sort label, are linked by arcs labeled with the names of grammatical features, where for some linguistic expression C, each feature f is taken to correspond to a property of C whose value is identified as the subtree originating at the node on which the arc labelled by f terminates. A graph corresponding to the example discussed in the preceding paragraphs is given in (1):

(1)

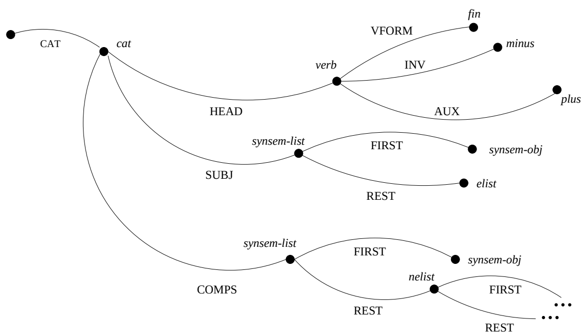

While such graphical representations seem to be intuitively more accessible than than the algebraic structures which essentially define feature strutures, they are still rather cumbersome to manipulate and far from perspicuous in displaying relevant information. These facts take on a particular urgency under the HPSG assumption that feature structures are complete objects—that is, total specifications of the grammatical properties of linguistic expressions. In these terms, the object in (1) is not a feature structure, but rather a small substructure; a full specification, as will become evident shortly, would involves the display of an enormous number of nodes and arcs. Hence, a different and far more perspicuous format is employed in providing descriptions of feature structures. The feature sub-structure in (1) satisfies the description given in (2):

(2)
$$\begin{bmatrix} \texttt{CAT} & \begin{bmatrix} \textit{cat} & \\ \texttt{HEAD} & \begin{bmatrix} \textit{verb} \\ \texttt{VFORM}\ \textit{fin} \\ \texttt{INV}\ - \\ \texttt{AUX}\ + \end{bmatrix} \\ \texttt{SUBJ} & \left\langle \textit{synsem-obj} \right\rangle \\ \texttt{COMPS} & \left\langle \textit{synsem-obj, synsem-obj, ...} \right\rangle \end{bmatrix} \end{bmatrix}$$

Such objects are called ATTRIBUTE-VALUE MATRICES, and are used both to provide descriptions of feature structures and to state constraints on them.

It often proves convenient to assume that sorts themselves are organized hierarchically, into networks of subsort/supersort relations. Every object that the grammar licenses must satisfy the signature declarations on all the sorts of which that object is a subsort, in addition to whatever other constraints apply to it beyond those in the signature (e.g., the Head Feature Principle discussed below). Schematically, one can depict such relationships among sorts in terms of a branching hierarchy. Each subsort is subject to a particular set of restrictions, imposed on it outside the signature properties themselves, in the constraint logic constituting the description language of the grammar. Elements lower (i.e, more specific) in the sort hierarchy inherit all restrictions imposed on the higher (i.e., more general) sorts dominating it. The objects in the HPSG model theory, the feature structures themselves, only contain maximally specific sort labels; but descriptions of (classes of) feature structures, including the constraints of the grammar, may refer to sorts of any level of inclusiveness.

For example, the very top of the sort hierarchy is, by convention, taken to be *object*, one of whose immediate subsorts is *sign*. Signs contain complete specifications for properties of classes of linguistic expression. The signature will identify one of the immediate subsorts of *sign* as *word* and the other as *phrase*; a further constraint will restrict the feature DAUGHTERS (DTRS) is to specification only for objects of the latter sort. The value of DTRS in turn must be of sort *constituent structure* (*cons-struc*), which has two subsorts: *coordinate-structure* and *headed-structure* (*headed-struc*), where the latter is subject to the restriction that it contain a specification for a feature HEAD DAUGHTER (HEAD-DTR). The sort *headed-struc* has the further subsorts *head-subject-structure, head-adjunct-structure, head-filler-structure*, and so on. Each of these subsorts is of sort *headed-struc* and so necessarily contains a specification HEAD-DTR; in addition, each may be subject to idiosyncratic restrictions. All phrases, for example, of sort *head-filler-struc*, where a filler is linked to a gap site, require a specification FILLER-DTR which is related to the feature specifications on the HEAD-DTR category in a specific fashion allowing the filler and gap site to share relevant properties, as discussed below.

It is frequently suggested informally by researchers within HPSG that such type hierarchies are cognitively realistic, insofar as they are crucial to characterizing psycholinguistic processes such as language acquisition. Although space limitations preclude detailed discussion, it should be clear that the constraint inheritance property of the type hierarchy sketched in this overview allows the actual lexicon itself to be relatively spartan, given a sufficiently rich hierarchy of sorts with associated constraints. This will be especially true if multiple lines of inheritance are permitted in the hierarchy. A given lexical item will have to bear only the feature information which is absolutely idiosyncratic; any information which can be predicted by the fact that it belongs to a certain class of lexical items can be built into a constraint on a sort to which those items belong, and if that class of items shares a certain property with a separate class of items, there will be a supersort under which the sorts of both respective classes appear, with the shared property identified as a constraint on the supersort. The sort hierarchy thus allows generalizations over the lexicon to be stated in a compact fashion, with fully specified lexical descriptions deducible in a straightforward monotonic fashion. Nonetheless, it is not evident at present to what degree the type hiearchy captures significant properties of grammars which are otherwise unstatable. In other words, whether the type hierarchy has more than an abbreviatory status is still an open question and depends, to a large extent, on what can be inferred from well-developed accounts of the interaction between formal grammar on the one hand and other linguistic phenomena, such as acquisition, processing or language change on the other.

2.2 The architecture of categories

The inventory of attributes or features in HPSG descriptions expressing properties of linguistic signs does not in itself determine the detailed structure of such signs. The latter is an empirical question and in the decade and a half in which work in HPSG has proceeded, there has been continuous rethinking of the specific feature geometry required to capture grammatical properties of natural language in a sufficiently restrictive way. Currently, for example, there is general agreement among HPSG theorists that lexical signs require specification for PHONOLOGY (PHON), SYNTAX-SEMANTICS (SYNSEM), MORPHOLOGY (MORPH), while phrasal signs require PHON, SYNSEM and DTRS attributes. The general organization of phrasal signs is given in (3), where descriptions such as (3) are to be understood as informal summaries of properties of the class of admitted feature structures:

(3)
$$
\begin{bmatrix}
\mathtt{PHON}\ \langle \ldots \rangle \\
\mathtt{SYNSEM} \begin{bmatrix} \mathtt{LOC} \begin{bmatrix} \mathtt{CAT} \begin{bmatrix} \mathtt{HEAD}\ \ldots \\ \mathtt{VAL}\ \ldots \\ \mathtt{ARG\text{-}ST}\ \ldots \end{bmatrix} \\ \mathtt{CONT}\ \ldots \\ \mathtt{CONX}\ \ldots \\ \mathtt{QSTORE}\ \{\ldots\} \end{bmatrix} \\ \mathtt{NONLOC}\ \ldots \end{bmatrix} \\
\mathtt{DTRS} \begin{bmatrix} \mathtt{HEAD\text{-}DTR}\ \ldots \\ \mathtt{NONHEAD\text{-}DTRS}\ \ldots \end{bmatrix}
\end{bmatrix}
$$

The value of SYNSEM is itself a set of specifications for LOC(al) features, reflecting locally relevant properties of the sign, and for NONLOC(al) features, encoding information which propagates over unboundedly large syntactic domains, such as information connecting fillers to gap sites, or information about *wh* formatives introducing information into relative or interrogative structures. LOC specifications identify syntactic properties of lexical or phrasal categories via the feature CAT, information relevant to logical aspects of interpretation via the feature CONT(ent), and contextual information via the feature CONX. LOC is also specified for a feature QSTORE which plays a significant role in tracking the scope of quantifiers, as described below. For nominal categories, an INDEX feature is defined as part of the CONT specification, which (ignoring a very small class of exceptions) encodes reference to individuals or events. Values for CAT are particularly highly structured; as briefly sketched earlier, this feature specifies valence information and also HEAD values, where the latter feature encodes properties of lexical items that are shared with all phrasal projections of those items, including part of speech, as well as other properties that depend on the part of speech of the item (e.g., case values for signs whose HEAD is of sort *noun*; auxiliary or nonauxiliary status for items of sort *verb*, and so on). Valence requirements are encoded in subspecifications of the feature VAL. The particular features which VAL is in turn specified for depends on the syntactic category of the head; thus, nominal lexical items select both complements (via a feature COMPS) and a specifier, identified by a feature SPR and manifest as a possessor, a determiner or nothing, while the VAL specification for verbs requires specification for a subject (encoded as a feature SUBJ) and COMPS, but not SPR.

In addition to the valence features SUBJ, COMPS, SPR, HPSG posits a feature ARG-ST which typically takes the form of a list-append operation carried out on the valence lists. In the case of verbs, for example, the COMPS list is appended to the singleton SUBJ list, so that the ARG-ST list is identical to the COMPS list with one extra element corresponding to the subject added at the top of the list. The ordering of the elements on the ARG-ST list is the same as that for the valence list, and expresses the notion of *relative obliqueness*, discussed below. The ARG-ST list is not a valence list, but rather a record of the verb's argument structure; it plays a major role in the nonconfigurational HPSG theory of coreference, outlined in §2.3.3.

This organization of features embodies certain well-motivated assumptions about linguistic modularity. For example, the value of each of the valence features is required by appropriate sort declarations to be a list of *synsem* specifications, where the feature geometry of the framework, as just noted, packages syntactic and semantic information about categories, but not phonological or subconstituent information, within *synsem* objects. The result is the exclusion of access by a selecting head to information about either the phonological form or the syntactic subconstituency of its arguments.

2.3 Capturing syntactic dependencies

The critical task facing any grammar is capturing systematic linkages in form between (possibly quite distant) substructures of a larger syntactic structure. Were it not for the existence of such linkages, or *syntactic dependencies* as they are generally referred to, natural language grammars would be far simpler than they actually are, since there would be little or no need to account for the sharing of often quite detailed information between different subcomponents of a sentence. Given the existence of such dependencies, a major goal of grammatical theory is the discovery of the optimal mechanism for expressing them.

The explicit specification of grammatical properties by means of highly structured descriptions, as sketched in the previous section, allows HPSG to capture syntactic dependencies economically. Agreement and case government are common examples of local dependencies, which typically hold between elements in a single clause, but other dependencies are more intricate, including several which have been taken as canonical demonstrations of the need for derivational relationships among syntactic objects. Such derivations typically serve to link the grammatical information shared among (possibly arbitrarily distant) structural positions. In HPSG, the linkage among the elements of the dependency is specified directly via the interaction of lexical properties with a small number of highly general declarative constraints, without the invocation of any mediating device such as movement.

2.3.1 Local dependencies

Local syntactic dependencies can in most cases be expressed by the correlation of properties of a selecting head with properties of a selected constituent, corresponding to some valence element. The simplest cases of this mechanism are syntactic selection and agreement. For example, the fact that the verb *give* must appear with an NP and a PP whose head is the preposition *to* is captured by identification of the verb's COMPS value as $\langle \boxed{1}\text{NP}, \boxed{2}\text{PP}[\textit{to}] \rangle$, where [*to*] abbreviates [PFORM *to*], identifying the morphological form of the preposition via a feature PFORM, and where the angled brackets conventionally denote a list whose elements are ordered from left to right according to their relative obliqueness, corresponding to a traditional hierarchy of grammatical relations. Direct objects, for example, are the least oblique elements on a verb's COMPS list. At the same time, the CONT value of *give* specifies a particular relation among an agent, a recipient and an object given. The HPSG approach to selection can be illustrated by the entry in (4), where we employ the following conventional notation:

- boxed alphanumeric symbols, called *tags*, are variables over nodes in the feature structures to which the descriptions containing those tags apply. Thus, when two features in a description appear with identical tags as their respective values, they have precisely the same node (and its associated subgraph) in the modeling object for those values;
- sorts, where relevant, are conventionally indicated in italics at the top of the description; and
- the notation $X_{\boxed{n}}$ indicates that the INDEX value of X is n.

Adopting these convention, we can provide the following partial lexical entry for *give*:

(4)
$$\begin{bmatrix} \text{PHON} & \langle \text{give} \rangle \\ \text{SYNSEM|LOC} & \begin{bmatrix} \text{CAT} & \begin{bmatrix} \text{HEAD} & \begin{bmatrix} \textit{verb} & \\ \text{AUX} & - \end{bmatrix} \\ \text{VAL} & \begin{bmatrix} \text{SUBJ} & \langle \boxed{1}\ \text{NP}_{\boxed{2}} \rangle \\ \text{COMPS} & \langle \boxed{3}\ \text{NP}_{\boxed{4}},\ \boxed{5}\ \text{PP}_{\boxed{6}}\left[\text{PFORM}\ \textit{to}\right] \rangle \end{bmatrix} \end{bmatrix} \\ \text{CONT} & \begin{bmatrix} \textit{give-relation} & \\ \text{AGENT} & \boxed{2} \\ \text{GIVEN} & \boxed{4} \\ \text{RECIPIENT} & \boxed{6} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Two general principles of the grammar are assumed here.

- The Valence Principle ensures that a phrase can be projected from a head just in case there is exactly one sister for that head corresponding to every *synsem* object on the head's COMPS list.

- The Head Feature Principle requires the HEAD specifications of a mother and its head daughter to be structure-shared (i.e., the arc from the mother category labeled HEAD and the arc from the head daughter category labeled HEAD terminate on the same node, in the graphic representation of admissible feature structures sketched earlier).

The interaction of these principles with the lexical entry given in (4) entails that the only feature structures that will be admitted are just those in which the valence of the head is satisfied. Thus, if we represent the various DTRS feature values in term of familiar branching tree structure, then the feature structure meeting the description in (5)a will be licensed, but those described by (5)b-d will not be:

(5) a.

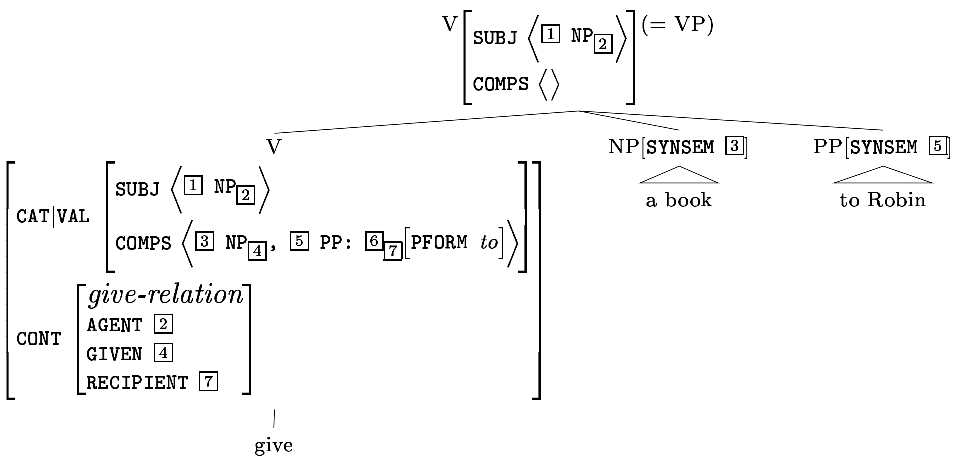

b.

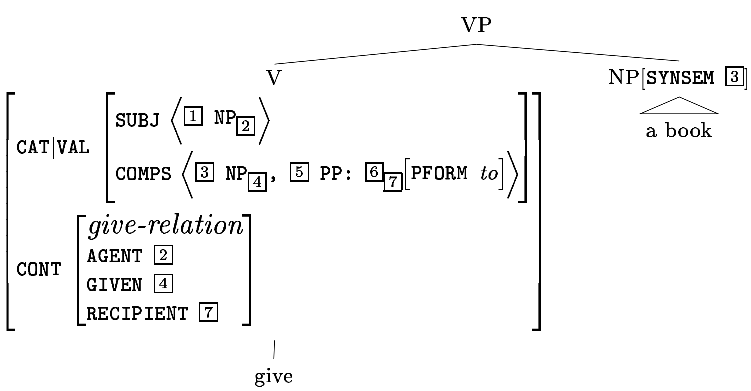

c.

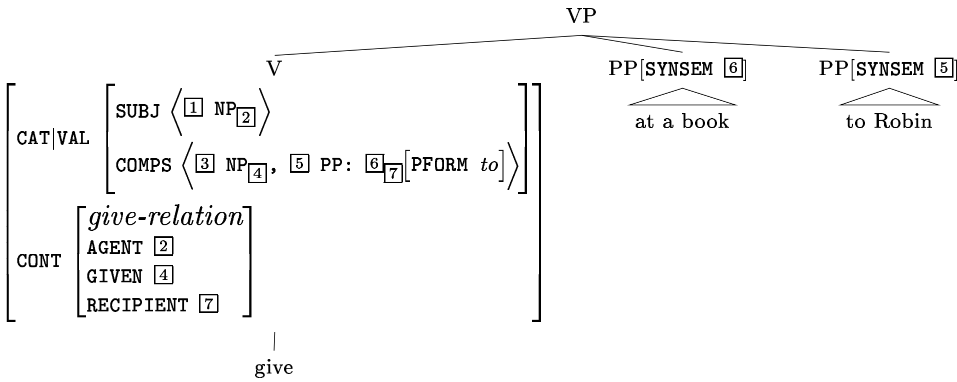

d.

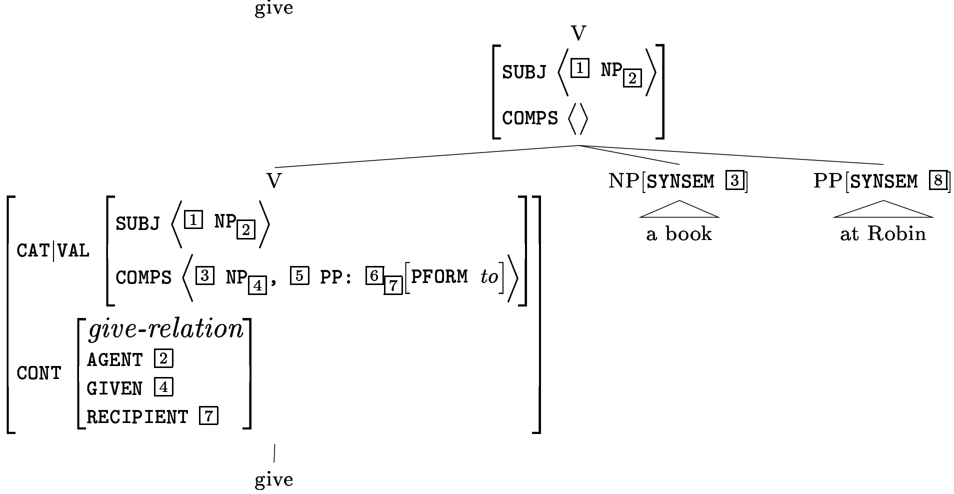

All four descriptions satisfy the Head Feature Principle, but only the first obeys the Valence Principle. Hence, only the VP depicted in (5)a is licensed by the grammar, yielding the relevant aspects of *give*'s subcategorization.

This same category selection mechanism can also enforce a correct match between the form of a verb and properties of its subject. Note, in the first place, that constraints on lexical types capture morphological generalizations; thus, it is unnecessary to separately stipulate that the third person singular form of *give* is *gives*, because a general constraint on the *third-sing* subsort of the *verb* sort automatically yields this form for a lexeme whose phonology is identified as [gɩv], and simultaneously imposes the specification [SUBJ ⟨ $NP_{\boxed{1}[\text{PERS } 3, \text{NUM sg}]}$ ⟩] on all verbs of this subsort. It then follows from the Valence Principle that the only NPs which will be able to combine with VPs projected from *gives* and morphologically parallel verbs are third person and singular.

Given the possibility just alluded to of constraining the subject selection properties of classes of verbal heads with reference to some other grammatical property of those heads, the standard local dependencies of English can be captured simply along the following lines:

- Consider the consequence of positing a lexical relationship which guarantees that for every verb with a [SUBJ ⟨ S ⟩, COMPS $\langle \alpha_1, \ldots, \alpha_n \rangle$] specification, there is a matching lexical entry with an identical phonology with the specification [SUBJ $\langle \mathrm{NP}_{it} \rangle$, COMPS $\langle \alpha_1, \ldots, \alpha_n, \mathrm{S} \rangle$], where *it* is a reserved INDEX specification uniquely associated with the dummy pronoun *it*. Nominals must in any case be subclassified to make expletives available for reference so that, for example, ambient weather predicates such as *rain* take only *it* as subjects. But the lexical relationship just sketched will interact with the Valence Principle to ensure that for every clause projected from verbs such as *bother or astonish*, as in *That Robin would even consider spying for an unfriendly power astonished us* and *For Dana to act like that constantly would bother most people*, there is a matching clause with an *it* subject and an identical set of complement daughters, along with an additional clausal daughter, e.g., *It astonished us that Robin would even consider spying for an unfriendly power*, *It would bother most people for Dana to act like that constantly.* Thus, extraposition is accounted for without having to posit a derivational origin for this construction.

- the active/passive correspondence can be treated as a lexical relationship in which, either by a lexical rule or a type constraint, active verbs are related to their passive counterparts in such a way that the COMPS list of the latter is that of the former minus the highest-ranking complement specification, where this 'lost' element reappears as the passive's SUBJ specification.
- The existential *there* construction can be captured by means of a partial lexical entry for the lexeme *be* as in (6):

(6)

$$\begin{bmatrix} \mathrm{PHON} \; \langle \mathsf{be} \rangle \\ \mathrm{SYNSEM|LOC} \begin{bmatrix} \mathrm{CAT} \begin{bmatrix} \mathrm{HEAD} \; \textit{verb} \\ \mathrm{VAL} \begin{bmatrix} \mathrm{SUBJ} \left\langle \mathrm{NP}_{\textit{there}}[\mathrm{NUM} \; \boxed{1}] \right\rangle \\ \mathrm{COMPS} \left\langle \boxed{2} \; \mathrm{NP}_{\boxed{3}\textit{ref}}[\mathrm{NUM} \; \boxed{1}], \begin{bmatrix} \mathrm{SUBJ} \; \langle \boxed{2} \rangle \\ \mathrm{COMPS} \; \langle \rangle \\ \mathrm{CONT} \; \boxed{4} \end{bmatrix} \right\rangle \end{bmatrix} \end{bmatrix} \\ \mathrm{CONT} \; \boxed{4} \end{bmatrix} \end{bmatrix}$$

where the subscript *there* notates a separate *index* subsort like *it*, and both are distinct from the subsort *ref(erential).* Under this lexical analysis, *There are two lions in the closet* will denote exactly what the predicate *in the closet* applied to *two lions* denotes, while the verb morphology is (indirectly) correlated with the person and number specifications of the post-copula NP.

- Properties of raising constructions (e.g., *Robin seems to be having a difficult time*) can be straightforwardly be captured by allowing raising verbs, such as *seem*, to bear specifications for the attribute SUBJ which are identical to those of their infinitival VP subjects; at the same time, the CONT specifications of raising verbs do not identify a semantic role for their subjects, only their infinitival VP complements. This simple lexical property immediately ensures that raising verbs preserve the subject-selection properties of their complements. It also yields other familiar properties of raising verbs, such as their preservation of meaning-invariance and idiomaticity under passivization. Similar observations follow for auxiliaries, if these are also taken to structure-share their subject values with their complements but not to associate these values with semantic roles.

2.3.2 Nonlocal dependencies

The selectional mechanism sketched in the previous section is not sufficient to guarantee the observed linkage between elements in the syntactic relationships usually referred to as unbounded dependencies. The difficulty can be illustrated in an example such as (7):

(7) a. $\left[_{\text{NP}} \text{Which} \left\{ \begin{array}{c} \text{busboys} \\ \text{*busboy} \end{array} \right\} \right]$ did Mary tell her Aunt Winona that Bruce wanted the headwaiter to ensure __were never seen filling her water glass again?

b. $\left[_{\text{NP}} \left\{ \begin{array}{c} \text{who} \\ \text{*whom} \end{array} \right\} \right]$ do you believe (that Robin thinks) __goes there often?

As these examples show, a significant amount of grammatical information is shared between the fronted constituent and the place from which some element is clearly missing, notated with an underline. The apparently displaced material in brackets is an NP, and the reader can verify that replacing it with an AP or a PP, or indeed any other phrasal category, gives rise to an ill-formed result. Furthermore, as the data in (7) attest, it is not only the gross categorial description that is shared between this filler and the gap site: information about the number required for subject/verb agreement is also part of the linkage. Case information is also evidently preserved; e.g., the requirement that subjects of a finite verb must be nominative is apparently imposed on arbitrarily distant filler associated with the gap. Such data were for decades taken as the strongest possible evidence that derivational machinery—by allowing the filler to satisfy contextual requirements at a level of structure prior to its displacement in subsequent stages—is essential for an adequate account of syntactic phenomena.

It is evident that the local mechanisms previously introduced cannot mediate such relationships. These mechanisms require the components of the dependency to be visible to a single lexical head, allowing properties of subjects and VPs to be correlated in appropriate fashion, giving rise to the dependencies observed, whereas in a nonlocal dependency, there is no single lexical head in whose VAL list specifications the properties of the filler and of the gap site may both be specified. Indeed, the filler in such constructions typically appears in a configurational relationship with a clause which by definition, is saturated for all valence features, so that there is literally no way for information about the filler to be linked by means of such valence features. The solution, inherited in HPSG from work by Gazdar (1981) and Gazdar et al. (1985), is to introduce a separate feature SLASH, whose values are a set of LOC specifications—all the information contained SYNSEM values except for information about nonlocal features, including SLASH itself— which is shared between a mother and at least one daughter. The propagation pathway for SLASH terminates in a gap under specific conditions ensuring that the crucial information about the filler noted earlier is preserved at the gap site. In HPSG, following the approach in GPSG, this is accomplished by defining a three-part mechanism for filler/gap linkages:

- The 'top' of such linkages is licensed by identifying a clause whose DTRS value of is of sort *head-filler-phrase*. This sort is constrained to specify a non-head daughter with a LOC value that is structure-shared with the SLASH value of the clausal head daughter. The effect is to guarantee the existence of structures of the form in (8):

(8)

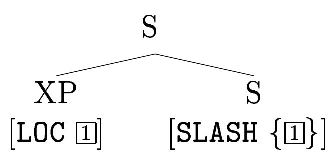

- The 'middle' of the linkage arises as the effect of a constraint, the Slash Inheritance Principle, which forces SLASH specifications on mothers to be shared with at least one daughter and vice versa, subject to conditions which ensure that the SLASH value at the top of the dependency is bound off at that point. Thus, once SLASH has been introduced into a structure, it must appear at every hierarchical level below the point of its introduction until it is terminated.

- HPSG theorists are divided on how termination of the SLASH pathway terminates. On one account, the SLASH specification is 'cashed out' by means of an empty category, while an alternative terminates the percolation of SLASH at the bottom of its path on a lexical head whose valence is reduced by an element corresponding exactly to the SLASH value of that head. For simplicity of exposition, the form of a filler/gap dependency is illustrated using the first approach:

(9)

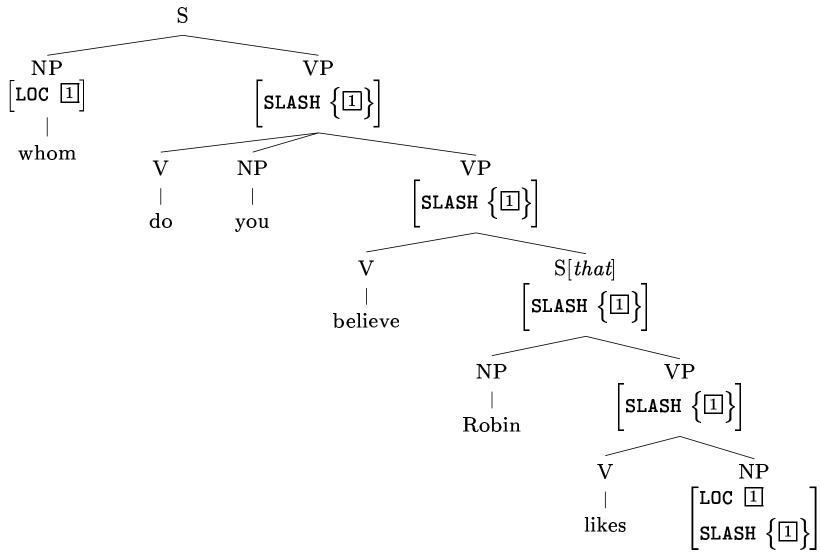

Given the existence of the lexical entry in (10), the SLASH path will terminate in a phonologically null category which—since it has a LOC value identical to that of the SLASH value (and of the filler)—preserves all the categorial and content information of the filler, and therefore must be compatible with any restriction imposed at the gap site.

(10)
$$\begin{bmatrix} \text{PHON} \ \langle\rangle \\ \text{SYNSEM} \begin{bmatrix} \text{LOC} \ \boxed{1} \\ \text{SLASH} \ \left\{\boxed{1}\right\} \end{bmatrix} \end{bmatrix}$$

2.3.3 Syntactic conditions on anaphora

The obliqueness ordering on valence elements alluded to above provides an essential component in the HPSG account of syntactic factors which enter into the determination of whether two given constituents can be interpreted as denoting the same entity. A very concise set of principles—the so-called *binding theory*—can provide an account of this relationship of coreference over a large range of linguistic phenomena. The binding theory defines, for each of a small number of types of referring expression, the conditions under which they may be coreferential with some other component of the linguistic sign in which they appear. The HPSG framework is perhaps unique among phrase structure-theoretic approaches in appealing to conditions on coreference which make no appeal to the configurational relationship between corefering elements. Rather, the basis of the binding theory is the obliqueness relationship defined on the ARG-ST list discussed above, in which all the elements selected by some lexical head are displayed in order of ascending obliqueness, with subjects taken to be the least oblique elements on such lists.

An argument of a given head is said to *locally o(bliqueness)-command* any more oblique coargument. More generally, it *o-commands* any more oblique coargument, the head of that coargument, any arguments of that head, and any other elements that the arguments of that head themselves o-command. For any lexical head, this recursive definition in effect picks out a path from some argument of that head through arbitrarily deep levels of structure linked by headship and selection. Its application to the statement of coreference possibilities is illustred below directly.

The HPSG binding theory can then be succinctly stated as follows:

A. Reflexive and reciprocal forms with referential local o-commanders must be coindexed by one such o-commander;
B. Pronouns which are not reflexive/reciprocal forms must not be coindexed with a local o-commander;
C. Non-pronominals cannot be coindexed with any other element.

The theory of coreference just stated provides a satisfactory account of a wide variety of coreference facts in English, illustrated in (11):

(11) a. Robin$_i$ admires him*(self)$_i$.
b. Pictures of him(self)$_i$ intrigue Robin$_i$.
c. It was only herself$_i$ that she$_i$ had to blame for what happened.
d. Robin$_i$ believes that people$_j$ admire her(*self)$_i$.
e. *She$_i$ believes that people$_j$ admire Robin$_i$.
f. *Who$_i$ does she$_i$ believe that people admire t$_i$?

In (11)a, the subject appears on the same ARG-ST list as the object, hence locally o-commands it. If the object is a reflexive, then by principle A it must be coindexed with the local o-commanding NP, and by the same token, if the object is a personal pronoun, it must *not* be bound by a locally o-commanding antecedent. In (11)b, on the other hand, there is no locally o-commanding antecedent, since the nominal head *pictures* has no other arguments ; hence the example vacuously satisfies principle A and reference is syntactically free, though constrained by pragmatic factors. By the same token, the absence of any local o-commanding antecedent in this example allows the personal pronoun to appear. Thus, the often-noted failure of reflexives and personal pronouns to exhibit completely complementary distribution follows straightforwardly from the HPSG binding theory. (11)c again illustrates how reflexives can appear freely in the absence of a local coindexed o-commander; here, *it*, though it o-commands the reflexive, does not have an index of sort *ref*; hence, by Principle A, it does not require a coindexed local o-commander. (11)d, however, a local, referential o-commander is present with which a reflexive fails to be coindexed, yielding ill-formedness. By the same token, however, a pronoun contraindexed with the local o-commander can appear, in accord with principle B.

Examples(11)e-f illustrate the ill-formedness resulting when nonpronominals are coindexed. In both cases, *she* as matrix clause subject is less oblique than the clausal complement of the matrix verb, and therefore (locally) o-commands this clause. In view of the preceding discussion, it therefore o-commands the head of this clause *admire*, and all the elements that *admire* itself selects, including *Robin* in e. and the trace of *who* (which is taken to be nonpronominal) in f. Hence these nonpronominal elements are o-commanded by, and coindexed with *she*, in violation of principle C, and (11)e-f are therefore ruled out as required.

2.4 Preliminaries to semantic interpretation

Every sign in an HPSG grammar is specified for a CONT value which provides the input to semantic interpretation. Nominal content values, of type *nom-obj*, identify an index and a set of restrictions on that index; thus, the noun *gift* optionally combines with two PPs, *of* $\mathrm{NP}_{\boxed{2}}$, *to* $\mathrm{NP}_{\boxed{3}}$, whose own CONT values are identified with those of their complement daughters, yielding nominals such as *gift of a book to Robin*, whose CONT has the form

(12)
$$\begin{bmatrix} \texttt{INDEX}\ \boxed{1} & \\ \texttt{RESTR(iction)} & \begin{bmatrix} \textit{gift-relation} \\ \texttt{INSTANCE}\ \boxed{1} \\ \texttt{GIVEN}\ \boxed{2} \\ \texttt{RECIPIENT}\ \boxed{3} \end{bmatrix} \end{bmatrix}$$

This partially saturated nominal will then combine with some determiner, such as *the* or *every*, whose own content identifies its quantificational properties, and whose lexical specifications allow it to take the content of a nominal head such as (12) as the restriction on the relevant quantifier. The semantic content of the quantifier is identified as the NP's QSTORE value, an implementation in a feature-based framework of earlier proposals made by Robin Cooper (see, e.g., Cooper 1983). QSTORE values of daughters are added to that of their mothers progressively up the tree, to the point where a quantifier is retrieved from the QSTORE of some category and added to the QUANTS value at that same node, thus fixing its scope. By the time the top of the tree is reached, all elements have been retrieved from QSTORE, with their left-to-right order identifying their relative scoping.

The preceding makes it clear that essentially two kinds of information must be tracked under this architecture: on the one hand, there is a kind of quantifier-free 'nuclear' content, and on the other the retrieval and ordering of scoping among quantifiers over structures of arbitrary depth. The former corresponds to the successive embedding of argument structures as parts of increasingly complex linguistic expressions, as imposed by the Valence Principle and other constraints of the grammar. The latter is expressible as the gathering of quantifier-contents into the QSTORE set at successively higher levels in phrasal structure and their removal at the point in the structure where they are required to take scope in order for a given reading to be represented. This dual percolation of information is regulated by the constraints already mentioned, and a separate condition, the Semantics Principle. The retrieval of quantifiers from QSTORE is nondeterministic, so that all possible scopings are realized, subject to whatever empirically motivated restrictions prove necessary.

3 Advantages and limitations

Among the various competing grammatical frameworks currently available, HPSG enjoys a number important strengths:

- The most prominent advantage of HPSG is one that it shares with other monostratal frameworks, viz., the ability to express compactly analyses of natural language phenomena incorporating significant generalizations without the assumption of derivational histories to encode the relevant properties of the phenomena described.
- The assumption of phrase structure configurations in natural language provides a straightforward basis for defining the percolation of information in syntactic objects. In particular, syntactic heads can serve as the locus of feature sharing between mother and daugher categories, a tendency evident in recent HPSG treatments of the distribution of nonlocal features and of QSTORE.
- Its formal foundations are mathematically explicit and well-understood, taking the form of a logic of constraint satisfaction which ensures a complete interpretation for all posited constraints. This property of HPS grammars guarantees that any inconsistent components of an HPSG analysis are in principle detectable.
- The monotonicity of HPSG's mechanisms for regulating the distribution of grammatical information in syntactic structures, combined with the formal explicitness already noted, makes computational implementation of HPSG analyses particularly straightforward. Indeed, much of HPSG's theoretical apparatus has developed in tandem with computational practice, e.g., the existence of efficient unification algorithms. As noted earlier, in certain quarters defaults have been introduced, but the nature of these default is essentially abbreviatory, in that every such HPS grammar can be compiled out to one in which no defaults appear, using more restricted formulations of the relevant constraints. The essentially abbreviatory status of such defaults ensures that, from a computational point of view, all current HPS grammars can be regarded as monotonic.
- The use of rich modeling objects such as typed feature structures allows the theory to combine a high degree of descriptive flexibility (conspicuous, for example, in the framework's ability to model various kinds of unbounded dependency construction) with a feature-geometric organization that captures important modularity properties of natural language grammars, as discussed earlier.

The downside of HPSG's phrase structure basis emerges in contexts where it is less clear that the elements involved in the relevant constructions are definable in highly configurational terms. Coordination, for example, manifests a number of phenomena in which it appears that string parallelism, rather than constituency, is the basis of a satisfactory account. HPSG currently lacks a detailed theory of coordination, and nonconstituent coordination, in particular, is problematic for phrase structure approaches. The greater combinatorial flexibility of categorial grammar probably compares most favorably to phrase structure approaches in this empirical domain. Currently, the lack of a well-developed approach to coordination is one of the most serious deficits in HPSG's theoretical coverage.

A second difficulty faced by current HPSG is a consequence of defining the relevant command relation for binding theory strictly in terms of valence. Since, under conventional assumptions about adjunct elements, adverbs and other modifiers are not lexically selected by heads, the prediction follows that the binding theory sketched above should be inapplicable to elements within adjuncts. But it is not at all clear that this prediction is empirically warranted. On the other hand, assuming that adjuncts are indeed lexically selected brings adjuncts into line with clearly selected elements, but again in a way that clashes with significant data. It is not clear, therefore, whether a strictly valence-based account of coreference possibilities is tenable, or if a configurational component must be reintroduced in order to bring the binding theory into line with the full range of relevant facts.

4 Future prospects

Despite the tendency of important discoveries to deflect established lines of research onto unforeseen pathways, certain trends already evident in HPSG are nonetheless very likely to become increasingly significant over the next several years. In particular, future work in HPSG will almost certainly focus on a number of issues where the current theory is notably incomplete.

- New approaches to problems posed by word order, in particular cases of so-called 'discontinuous constituency', are likely be be developed based on what has been called the *linearization framework* within HPSG, in which word order facts are associated with a descriptive domain related to, but formally separate from, that which defines constituency.
- The 'constructional' approach to syntactic phenomena pioneered in Sag 1997, relying on elaboration of the phrasal sort hierarchy to capture both regular and idiosyncratic syntactic properties, will become increasingly prominent. Exceptional behavior, e.g., the fact that relative clause modifiers with neither *wh* nor *that* formatives cannot be attached to nominal heads corresponding to their subjects (*the guy Robin criticized is here/*the guy criticized Robin is here*), or the fact that PPs containing *wh* can link infinitival relative clauses to nominal heads, but that *wh* NPs themselves cannot (*something about which to complain/*something which to complain about*) will increasingly be handled through sort declarations which encode them as construction-specific constraints. It must be stressed, however, that there is far from unanimous acceptance of this 'constructional turn'.

- The semantic component of HPSG is still seriously underdeveloped, following its dissociation from the situation-theoretic perspective which guided much earlier work. As things stand, the CONT and QSTORE features provide information which constitutes a kind of logical form representation, but there currently no generally accepted constraining relationship between such representations and empirically motivated formal models. A particularly stubborn obstacle to the formulation of such relationships is the difficulty posed by major outstanding empirical problems, such as the semantic nonequivalence in opaque contexts of intensionally equivalent objects. Research currently in progress in HPSG is aimed at eliminating this deficit of the theory, and will probably intensify in the near future.
- Finally, there has already been work on constraint-based phonology which extends the HPSG approach of using sorted feature structures to state mutually constraining representations to phonology, with the ultimate objective of providing a completely nonderivational theory. An overview of work along these lines, with pointers to relevant sources, is given in Bird and Klein 1994, and the application of sorted feature structures and other aspects of the HPSG formalism will probably become increasingly important for foundational work in phonology.

Article Headword Definition:

Head-Driven Phrase Structure Grammar: a monostratal theory of natural language grammar, based on richly specified lexical descriptions which combine according to a small set of abstract combinatory principles stated as formulae in a constraint logic regulating, for the most part, the satisfaction of valence and other properties of syntactic heads. These constraints, applying locally, determine the flow of information, encoded as feature specifications, through arbitrarily complex syntactic representations, and capture all syntactic dependencies—both local and non-local—in elegant and compact form requiring no derivational apparatus.

Glossary Items:

monostratal: refers to a framework for licensing linguistic expressions whose licensing conditions are all defined on a single level of representation, as opposed to derivational frameworks in which for all linguistic expressions, licensing takes the forms of a succession of representations whose terminal stage provides the structural representation of that expression.

feature structure: a mathematical object which is taken in HPSG to model the properties of some linguistic expression, comprising a set of sorted nodes U and a set of unary functions from U to U, where each function in the set is represented as a feature (or attribute) label. An HPSG grammar of some language L is a complex admissibility requirement on the set of feature structures whose membership corresponds in one-to-one fashion to the linguistic expressions of L.

type/sort: an equivalence class in a partition of entities in of some domain, typically on the basis of shared properties which can be expressed as constraints on every member of the type.

signature: formally, a triple $\langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle$ where $\mathcal{S}$ is a set of sort names and $\mathcal{F}$ a set of feature names, and an assignment function $\mathcal{A} \in \wp(\mathcal{S})^{\mathcal{S}^{\mathcal{F}}}$, i.e., an explicit mapping from the set of functions to the set of mappings from sorts (for which the feature name is defined) to the powerset of sorts (the set of things that the feature can have as a value). Informally, the signature identifies the features which some description of a given sort must be specified for and the kinds of things that a feature specified for some sort can return as possible values.

attribute/value matrix (AVM): a descriptive device employed in HPSG to display the grammatical properties of the class of feature structures meetiing that description, consisting of a column of feature names (usually written in upper-case) on the left and the values—possibly complex values taking the form of AVMs themeselves—on the right.

relative obliqueness: a ranking among dependents of a head in terms of a notion of relative prominence or weight of the grammatical relations that they respectively bear to the head, with such prominence typically manifest in terms of, e.g., accessibility to extraction, and inversely correlated with degree of obliqueness; for verbs, for example, subjects are the least oblique (hence highest ranked) dependent, with direct objects next highest, and so on.

binding theory: a theory of the syntactic conditions regulating coindexation possibilities holding among anaphoric elements and their antecedents.

local o-command: a relationship that holds between two elements α and β on an ARG-ST list just in case α (i) is less oblique than β (i.e, outranks β) or (ii) locally o-commands some element γ such that β is a valent of γ.

o-command: a relationship that holds between two elements α and β just in case (i) α locally o-commands β; (ii) β is the head of some γ such that α o-commands γ; β is a valent of some γ such that α o-commands γ.

Further reading

Borsley, Robert. 1999. *Syntactic Theory* (2nd edition). London: Arnold.

Borsley, Robert. 1996. *Modern Phrase Structure Grammar.* Oxford: Blackwells.

Green, Georgia and Robert D. Levine. 1999. Introduction. In *Studies in Contemporary Phrase Structure Grammar*, ed. by Robert D. Levine and Georgia Green. Cambridge: Cambridge University Press.

Hukari, Thomas and Robert D. Levine. 1996. Phrase structure grammar: the next generation. *Journal of Linguistics* 32.465–496.

Kathol, Andreas. 2000. *Linear Syntax.* Oxford: Oxford University Press.

Pollard, Carl and Ivan Sag. 1987. *Information-based Syntax and Semantics.* Stanford: CSLI.

Pollard, Carl and Ivan Sag. 1994. *Head-driven Phrase Structure Grammar.* Chicago: University of Chicago Press.

Sag, Ivan and Thomas Wasow. 1999. *Syntactic Theory: a Formal Introduction.* Stanford: CSLI.

References

Bird, Stephen, and Ewan Klein. 1994. Phonological features in typed feature systems. *Computational Linguistics* 20:455–491.

Brame, Michael. 1976. *Conjectures and Refutations in Syntax and Semantics.* New York: North-Holland.

Cooper, Robin. 1983. *Quantification and Syntactic Theory.* Dordrecht: Reidel.

Culicover, Peter W., and Wendy Wilkins. 1984. *Locality in Linguistic Theory.* New York: Academic Press.

Gazdar, Gerald. 1981. Unbounded dependencies and coordinate structure. *Linguistic Inquiry* 12:155–84.

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

Gazdar, Gerald, Geoffrey K. Pullum, Robert Carpenter, Ewan Klein, Thomas E. Hukari, and Robert D. Levine. 1988. Category Structures. *Compuational Linguistics* 14:1–19.

Pollard, Carl J., and Ivan Sag. 1987. *Information–Based Syntax and Semantics.* Stanford University: Center for the Study of Language and Information.

Pollard, Carl J., and Ivan Sag. 1994. *Head–Driven Phrase Structure Grammar.* Chicago: University of Chicago Press.

Richter, Frank, Manfred Sailer, and Gerald Penn. 1999. A formal interpretation of relations and quantification in HPSG. In *Constraints and Resources in Natural Language Syntax and Semantics.* 281–298. Stanford: Center for the Study of Language and Information.

Sag, Ivan. 1997. English relative clause constructions. *Journal of Linguistics* 33:431–484.