

**SECOND LANGUAGE PRODUCTION: COPING WITH
GAPS IN LINGUISTIC KNOWLEDGE**

Michael Gasser

**April 1986
CSD-860082**

SECOND LANGUAGE PRODUCTION: COPING WITH GAPS IN LINGUISTIC KNOWLEDGE

Michael Gasser
Applied Linguistics Program
and
Artificial Intelligence Laboratory
UCLA

1. Introduction and Motivation

This paper reports on a portion of a project to develop a model of language production in an adult Japanese learner of English as a second language. The model is being incorporated into a computer program called CHIE.*

What does a model of language production have to offer second language acquisition theory and second language pedagogy?

1. Increasingly language production and comprehension are seen as aspects of language rather than as non-linguistic processes that operate on the linguistic system (e.g., Winograd, 1983). Language acquisition, then, is in part the acquisition of the ability to produce language.† Obviously a prerequisite for a theory of language acquisition is an understanding of what it is that is acquired; the CHIE project aims to characterize the production end of language at a point in the process of acquisition.
2. Learners, in particular adults, are often faced with expressing notions or satisfying communicative goals for which they are linguistically unprepared. In such circumstances it may often be untenable to simply abandon the attempt to communicate what they intended to. The alternative is to make creative use of whatever knowledge they have, of knowledge of the first and second languages as well as of the world, to attempt to get their point across. When a similar circumstance arises again, the same strategy is likely to be applied. Over time this may lead to the development of a relatively fluent pidgin which is established enough to be more or less "fossilized" (Selinker, 1972), that is, no longer amenable to instruction. (For the relationship between second language acquisition and the development of pidgins, see Andersen, 1983.) An understanding of how this process of pidginization takes place is crucial in second language acquisition research; it would appear to be rooted in learners' responses to communicative demands during language production.

The research reported on here was supported in part by a grant from the ITA Foundation.

*Chie is a Japanese woman's name and the Japanese word for 'wisdom'. The program is written in T, a dialect of LISP, and runs on Apollo Domain computers.

†Even if the Chomskyan performance-competence distinction is maintained, a linguistic theory must be supplemented with work on language comprehension and production because the theory is invalid unless it is shown that linguistic knowledge as it is characterized can be used to produce real utterances in context.

3. Most language acquisition is generally assumed to take place during comprehension, and one influential model of second language acquisition (Krashen, 1981) attributes essentially all of acquisition to what the learner does with linguistic input. However, there are several ways in which acquisition may also occur during language production.
 - a. When a particular form has just been associated with a function or meaning during comprehension, we would expect it to be weak, i.e., relatively inaccessible. Under these circumstances it might be possible to make use of the association during production only when there are few other processing demands placed on the speaker. Accessing the association would have the effect of strengthening it, that is, making it more accessible for later processing. A model of language production in a learner could explicate the nature of "processing demands" and reveal under what circumstances the learner can use and strengthen new form-function relationships.
 - b. As noted in 2. above, learners often discover gaps in their knowledge of the second language during production. With no help from outside these circumstances may eventually lead to the development of pidginized speech. If, however, the learner is provided with a rule or form appropriate to the context, the situation can result in acquisition. If the current emphasis on contextualized, communicative, meaningful practice in language pedagogy (e.g., Brumfit & Johnson, 1979; Hatch et al., in press; Nattinger, 1984) is well-founded, then we would expect instruction of this type to be superior to that not occurring in the context of genuine language production. A model of second language production could thus provide a basis for a principled theory of language instruction.
4. Most of the data which second language researchers have to work with are the output of a language production system. Before the data can be interpreted, the processes through which they arose need to be understood. One area that has been of interest is that of the avoidance of second-language structures which have not been fully learned (Schachter, 1974; Madden et al., 1978). A model of the production of a second language should help to clarify under what conditions avoidance takes place.

What takes place during the production of language, then, is of great importance in the acquisition process. In particular, a model of language production in a learner needs to address the issue of how the learner copes with gaps in linguistic knowledge. This paper is a first step towards a computational account of the strategies employed in such situations.

2. Second Language Production Errors

Speech errors made by second language learners may be due either to faulty associations of concepts with linguistic patterns or to the lack of associations between concepts and patterns. The former are a reflection of learning processes and the latter a reflection of production processes, but as I have argued, the production processes may result in learning. In some cases, it may be possible to interpret an error in either way, that is, as a result of a non-native-like association that has already been learned or as the result of an production strategy applied when there is no suitable association in memory. In this paper I deal with only production-related errors, those resulting from a gap in the learner's knowledge of the target language.

Gaps in linguistic knowledge result in two sorts of errors. On the one hand, there are errors with grammatical morphemes which express largely redundant meanings in the target language. Speakers usually seem to be unaware of such errors, and thus they need have no strategies at all to deal with the gaps that result in these errors. There are two of them in the following utterance; note that they do not interfere significantly with communication.

- (1) He wonder because there are only two basket.*

On the other hand, there are errors which are potentially more serious because a misunderstanding may result. In such cases the speaker is more likely to be aware of the problem, and processing may be noticeably interrupted. The problem may be lexical as in the next two examples. In (2) the speaker is trying to describe a scene in which two people are leaving each other after an encounter.

- (2) And uh passing away each other pass away
 no uh passing passed
 [Native speaker: Left each other?]
 Left each other.

She seems to be aware that *pass away* is not the most appropriate expression here but cannot find a better one. In the following example the lexical item is selected more readily.

- (3) I don't like this tea because it's too thick (= strong).
 (Takahashi, 1985)

Note that it is impossible to determine whether this error results from previous learning or a selection made during production in the context of a gap in linguistic knowledge. (dg In either case, this is a clear example of first language transfer; the Japanese adjective *koi* when used to describe a solution has the general sense 'concentrated', corresponding to both English *thick* and *strong*.)

In other cases where a gap is perceived by the speaker, the problem is the unavailability of an appropriate syntactic pattern to express a particular relationship. The sentences in (4) illustrate such a problem with the English relative clause pattern. The speaker has read a story in Japanese, her first language, and is retelling it in English and in Japanese. The story concerns two men who have found a hatchet lying on the ground. At this point the narrator wants to assert the fact that the man who had lost the hatchet was searching for it. In relating the story in Japanese, she uses a relative clause to distinguish this man from the other two:

- (4a) Soshite sono teono o otoshita otoko ga
 then that hatchet OBJ dropped man SUBJ
 'Then the man who dropped the hatchet

 sono teono o sagashite imashita.
 that hatchet OBJ was seeking
 was looking for the hatchet.'

*Unless otherwise indicated, examples were collected by the author.

†McShane & Dockrell (1983) note a similar difficulty in interpreting output data for first language acquisition.

In the English version, however, she seems to be unable to use a relative clause to perform this function and instead produces two sentences:

(4b) And then that saw is somebody's.
Somebody's looking for its saw.

There are, of course, problems with these sentences other than the one which interests us here: *saw* is used for 'hatchet' and *its* for 'his', and *somebody* in the second sentence incorrectly indicates a definite reference. To facilitate comparison with (4a), then, it will help to correct these other errors and recast (4b) in the following form:

(4b') And the hatchet was somebody's.
(Somebody had dropped the hatchet.)
He was looking for his hatchet.

The effect of the two sentences (assertions) is to imply to the listener that s/he cannot already infer the existence of the owner of the hatchet or of the act of losing it. In this context this is clearly not what the speakers intended.

In this paper I will show how both the distinction between errors such as those in (1) and errors such as those in (2), (3), and (4) and the particular behavior illustrated in (3) and (4b) make sense in terms of a model of language production as the selection of patterns associated with input goals.

3. Representation

Any model of language processing must deal with representation of

1. concepts, that is, general knowledge of the world (semantic memory) and knowledge of specific states, events and entities (episodic memory)
2. linguistic form, and
3. associations between concepts and form.

In CHIE there has been an attempt to represent all knowledge in a uniform fashion, in the form of a semantic network (see Flickinger et al., 1985 for a similar approach). This is in the spirit of the "unified" (e.g., Anderson, 1983) as opposed to the "modular" (e.g., Fodor, 1983) view of cognition. The next section introduces the general structure of memory in CHIE, section 3.2 discusses the organization of world knowledge, section 3.3 describes the representation of linguistic form, and section 3.4 looks at communicative goals, linguistic planning, the associations between concepts and forms, and the process of language production.

3.1. General Memory Organization

Both world knowledge and linguistic knowledge in CHIE are organized in a semantic network. A semantic network is a knowledge representation formalism in which nodes represent concepts and links the relationships between concepts (see Fahlman, 1979 and Brachman & Schmolze, 1985 for examples of semantic networks used in artificial intelligence and Anderson, 1983 for an example of a semantic network in a psychological model). Concepts include entities, states, and events. For CHIE I assume that a concept *is* the node that represents it, and I shall therefore use the terms **concept** and **node** more or less interchangeably. In CHIE, as in many semantic networks, nodes are organized into units called **descriptions**.^{*} A description functions as the definition of a concept referred to as the **head** of the description. Each description contains one or more links, called **is-a** links, from its head node to the **parents** of the head node, that is, the classes which it is a member of. For example, in the description of TEA[†] there are **is-a** links to the concepts SOLUTION and CONSUMABLE-SUBSTANCE, and in the description of INGEST there is an **is-a** link to the concept ACT. SOLUTION is a parent of TEA and TEA a **child** of SOLUTION. In addition, a description usually has one or more **roles**, concepts that are seen as “belonging” to the head concept. Roles are connected to the head by **has-a** links. For example, in the description of the concept SOLUTION there is a **has-a** link from the SOLUTION node to a CONCENTRATION role, and in the INGEST description there is a **has-a** link to an ACTOR role. The significance of these links is that every solution has an associated concentration and every act of ingesting has an associated actor. Roles are concepts in their own right and can have their own descriptions. For example, the ACTOR role in the INGEST description has an **is-a** link to the node for ANIMAL. This link signifies that only animals can ingest things (CHIE knows nothing about Venus flytraps).

The **is-a** links connecting concepts give the network a hierarchical structure. At the bottom of the hierarchy are nodes representing **individual** concepts; concepts higher up are **types**, that is, representations of typical instances. I shall also refer to individuals as instances of types. Nodes such as TEA and INGEST are types; these represent CHIE’s general knowledge about tea and acts of ingesting. Instances of these concepts, for example, the particular tea referred to in (3) above, are individuals. Names of individuals consist of the name of a type with an appended number, e.g., TEA2. Two concepts may be ordered with respect to their place in the hierarchy: if there is a path of **is-a** links leading upward from node A to node B, then B is said to be an **ancestor** of A and A a **descendant** of B.[‡]

When roles are specified in a description, there is always some restriction on the values they can take. They may be members of a particular class (may be assigned to parent concepts) or may be restricted to (equivalent to) a given individual concept. For example, as we saw above, the ACTOR role of INGEST is restricted to ANIMALS, and in a particular instance of INGESTing, say, the one referred to by the sentence *Chie drank some green tea*, the ACTOR is equivalent to a particular ANIMAL, identified as “Chie”. For the latter case we need a third type of link, an **equivalence** link, to indicate that two nodes represent the same entity. In this example there is an equivalence link joining the ACTOR role of the individual act (call it INGEST1) and the node for CHIE. The roles and role restrictions specified in a type description are all implicitly **inherited** by the descendants of the type. Thus there is no need to specify a

^{*}In most cases the terminology used in CHIE is borrowed from Fahlman (1979).

[†]Names of concepts, which are assigned only for convenience and have no special significance for CHIE, are written in small capitals.

[‡]The ordering is only a partial one, however; the concepts MAN and POTATO, for example, cannot be ordered in this way.

CONCENTRATION role for TEA because this is already specified in the SOLUTION description, and, since TEA is a descendant of SOLUTION, it automatically has a CONCENTRATION. What this means is that a description for a concept need only include the roles and/or role restrictions that distinguish that concept from its ancestors. Inheritance is normally implicit, but it can be made explicit when needed. Below we shall see how explicit inheritance is used in spelling out the form of individual phrases on the basis of knowledge about phrase types.

Figure 1 illustrates some of the notions discussed in this section as well as the graphical notation which I will make use of in this paper. The following notational conventions are used:

1. Rectangles represent THINGS, hexagons STATES, and octagons EVENTS (including ACTS).
2. Nodes with heavy borders are description heads; others are roles.
3. Nodes with inverted labels are type nodes; others are individuals.
4. Nodes with dashed borders are inherited roles.
5. Has-a links are simple line segments with the role always shown to the right of its head.
6. Is-a links are single-headed arrows and equivalence links double-headed arrows.
7. Dotted links join inherited roles to corresponding roles in descriptions of ancestors.

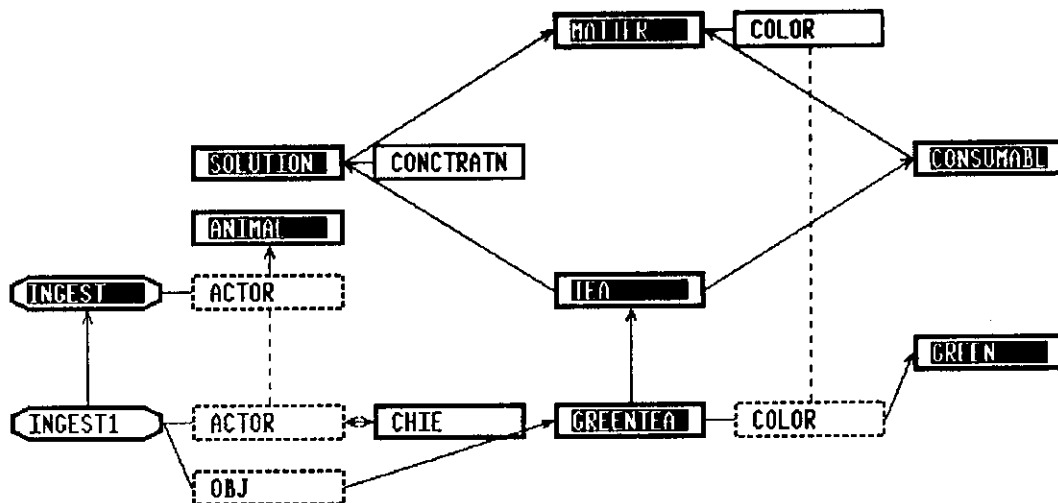


Figure 1

Note that in Figure 1, and in all figures in this paper, most descriptions are not shown completely. For example, no role specifications at all are given for TEA and CONSUMABLE.

3.2. World Knowledge

CHIE contains a minimal store of general semantic knowledge, for example, the knowledge that SEEing something implies KNOWing its LOCATION; that LOSing something implies not KNOWing its LOCATION; and that in order to gain CONTROL of something, one normally needs to GO to its LOCATION. CHIE also has an episodic memory for a small set of facts. This memory consists of suitably linked instantiations of concepts in semantic memory. It is assumed that episodic memory is non-linguistic and that, while each person may represent certain aspects of the world in an idiosyncratic way, all people, including those from different cultural backgrounds, share a large core of basic world knowledge. In this study cultural differences are not treated. More specifically, then, the assumption is that the facts in CHIE's episodic memory would be encoded in essentially the same form by a native speaker of English or a native speaker of Japanese. In order to allow for either Japanese or English to be generated from the same memory, individual concepts in episodic memory need to be instantiations of relatively primitive concept types (Schank & Abelson, 1977). For example the ACT described by the sentence *the man looked for the hatchet* is represented in a way which could be paraphrased as follows: MAN3 USED-A-PLAN of ATTENDING his EYES in some unspecified direction to achieve a GOAL of KNOWing the LOCATION of HATCHET1.

3.3. Patterns

Decisions about how knowledge of linguistic form is to be organized need to be based on the ways in which that knowledge is used during comprehension and production. In particular, this means that the units of linguistic analysis are those that can be associated with particular meanings (i.e., concepts) or functions (i.e., communicative goals). These units will be referred to as **patterns**. Patterns are clearly of varying generality; we want them to include relatively fixed phrases as well as syntactic structures. At the most specific end are items such as *that's the way the cookie crumbles* and at the most general, items such as **noun phrase with relative clause**. Patterns of course have an internal structure. Each of the constituents in a pattern may be specified precisely, that is, it may be restricted to a particular morpheme; it may be specified as belonging to a particular class of morphemes, for example, **noun**; it may be a pattern in itself; or it may be an instance of the SPEECH-ACT type (discussed in 3.4.). Constituent ordering is also a part of most patterns. For ordering it seems to be sufficient to have the following possibilities for a constituent: the stipulation that it precede or follow (though not necessarily immediately) some other constituent or the stipulation that it be pattern-initial or pattern-final. This explicit specification of constituent ordering is similar to the ID/LP rule format of Generalized Phrase Structure Grammar (Gazdar et al., 1985).

With the addition of sequence links, all of this fits into the semantic network scheme used to represent conceptual memory. In CHIE patterns are arranged in a hierarchy with lexical patterns lower and syntactic patterns higher in the hierarchy. The network approach permits the integration of two important trends in the representation of linguistic knowledge: that of phrase structure and functional grammars in linguistics (Bresnan & Kaplan, 1983; Flickinger et al., 1985; Gazdar et al., 1985; Kay, 1979) and psycholinguistics (Pinker, 1984; Stemberger, 1982) and that of phrasal lexicons in natural language processing work (Becker, 1975; Jacobs, 1985; Wilensky & Arens, 1980) and also to some extent in linguistics and psycholinguistics (Fillmore, 1979; Herskovits, 1985; Pawley & Syder, 1983). It also shares many features with connectionist approaches to language processing (e.g., Dell, 1985), which appear to be promising psychological models.

Winograd (1983, pp. 327-328) has discussed the advantages of the functional and phrase structure grammars. For our purposes the most important of these are that they allow for direct correspondences between patterns and meanings and that they are nondirectional and thus usable in both comprehension and production. Consider the NP with relative clause pattern, which in a phrase structure grammar is simply one possible realization (in CHIE one of the children) of the NP pattern. It is a fairly straightforward matter to associate with this pattern (or "rule" in a more conventional phrase structure approach) a specification of its semantics: the fact referred to by the relative clause is being attributed to the entity referred to by the NP; in other words, the entity fills one of the argument roles in the fact. At the same time, these semantic relations can be read in either of the two possible directions corresponding to what happens in production and comprehension. In production the need to attribute the fact to the referent would be realized as the complex NP, while in comprehension the structure would cause the listener/reader to create or search for an entity with the appropriate relationship to the fact. Similarly, the basic syntactic fact expressed by the phrase structure rule itself is interpretable in either direction. In production a slot for an NP of any type could be realized as the complex NP, while in comprehension the complex structure would be recognized as an NP.

The main idea behind phrasal lexicons is that a good deal of syntactic and morphological information can be built into lexical patterns of varying generality (referred to as "structural formulas" by Fillmore (1979) and "lexicalized sentence stems" by Pawley & Syder (1983)). This obviates the need to resort constantly to the syntactic level, which, as Pawley & Syder (1983) argue, would greatly reduce fluency. An example given by Pawley & Syder is the pattern *who (the EXPLETIVE) do-PRES NP_i think PRO_i be-PRES?*. In using this pattern to produce the sentence *who the hell does that guy think he is?*, a speaker avoids much of the syntactic work that is normally thought to be involved in formulating such a sentence. "Do-support" and "WH-fronting", for example, are already embodied in the pattern. In a language acquisition context one would expect the lexicon to contain an especially large proportion of such patterns (Hatch, 1983, p. 100).

In CHIE's hierarchy of patterns, descriptions of pattern types correspond to phrase structure rules, with the constituents as roles of the patterns. The distinction between role and head in CHIE's patterns corresponds to some extent to the distinction in Kempen & Hoenkamp's Incremental Procedural Grammar (1982) between **category procedures** (e.g., S, NP) and **functional procedures** (e.g., subject, NP modifier). One difference is that these categories are often collapsed in CHIE because a node can be both a role within a description and the head of a description at the same time. Lexical patterns have some or all of their roles (constituents) specified (i.e., linked to particular morphemes) and are thus relatively low in the pattern hierarchy. Particular sentences (phrases, etc.) are viewed as individual patterns which are instances of pattern types, from which they inherit their constituent roles. All individual patterns other than pro-forms have a lexical parent pattern from which they inherit one or more lexical morphemes and at least some of their syntax and usually at least one syntactic parent pattern from which they inherit their remaining syntax and possibly grammatical morphology. Figure 2 illustrates how the pattern *this tea* would be represented. Sequence links are indicated by arrows with filled heads. Names of morphemes are preceded by a slash, names of lexical patterns by a vertical bar. Thus \uparrow TEA indicates a NOMINAL with the morpheme /TEA as its SYNTACTIC-HEAD. The noun phrase NOM1 inherits its SYNTACTIC-HEAD role from the MAJOR-PATTERN type and its DETERMINER role from the DETERMINER-NOMINAL type. (Roles of the type concepts are not shown in the figure.) The DETERMINER-NOMINAL type also specifies the ordering of the constituents. The specification of the SYNTACTIC-HEAD role as equivalent to the morpheme /TEA is inherited from the \uparrow TEA type, and the specification of the DETERMINER role as equivalent to /THIS is inherited from the \uparrow THIS type. The assignment of NOM1 to the syntactic pattern NP stipulates its realization as a noun phrase rather than a clause of the form *X is tea*, which would also be an instance of the lexical pattern




Figure 2

E/TEA. Note that patterns are represented as EVENTS (ACTS, in fact) rather than THINGS. The reason for this will become apparent in the next section. Note also that there is no semantics associated with NOM1; it is a representation of structure only. Semantics is dealt with in 3.4.

CHIE of course requires both English and Japanese patterns in its repertory. For syntactic patterns these are represented as children of a higher-level language-independent type. For example, under the CLAUSE pattern type there is a JAPANESE-CLAUSE and an ENGLISH-CLAUSE.

3.4. Language Production

With this view of patterns, the task in language production becomes one of the selection of a set of patterns and their combination to produce a surface utterance. The input to this selection process consists of communicative goals. These include the desire to have the addressee

1. add a concept or a link to memory
2. locate a known concept in memory
3. perform a particular act
4. provide information to the speaker.

Because there is no straightforward one-to-one correspondence between goal types and pattern types, a planning mechanism is required to perform language production (Appelt, 1982). All planning (Schank & Abelson, 1977; Wilensky, 1983) starts with a goal of some sort, which is a state that the planner wishes to bring about. The planner attempts to find a plan, a kind of act, which can satisfy the goal. Often there is a canned plan (a "normal plan" in Wilensky's terms) associated with a particular goal type; in other cases, general strategies ("meta-plans" in Wilensky's terms) are used to generate an alternative plan. Once a

plan has been proposed, it may turn out that it has certain preconditions; these lead to further goals and recursive planning. There is also the possibility that a completely satisfactory plan cannot be found or generated and that the planner will have to settle for partial goal satisfaction.

In CHIE planning begins with the creation of an instance of the concept INTEND. INTEND is a high-level type which has roles for a PLANNER, a GOAL, and a PLAN which the PLANNER hopes will achieve the GOAL. The instance of INTEND that is created has its PLANNER and GOAL specified, and the task is to fill in the PLAN, which is an instance of some type of ACT. This is a matter of locating the information which the system has about how particular types of GOALS are associated with particular types of PLANS, that is, about the "normal" PLANS for GOALS. This can be seen as a process of classification: INTEND has subtypes which specify PLAN types for GOAL types, and the planning mechanism classifies the INTEND instance by attempting to match its GOAL with the GOAL role of one or more of these subtypes. Two descendants of INTEND which are relevant for this paper, ASSERT and DEFINITE-REFER, are shown in Figure 3.

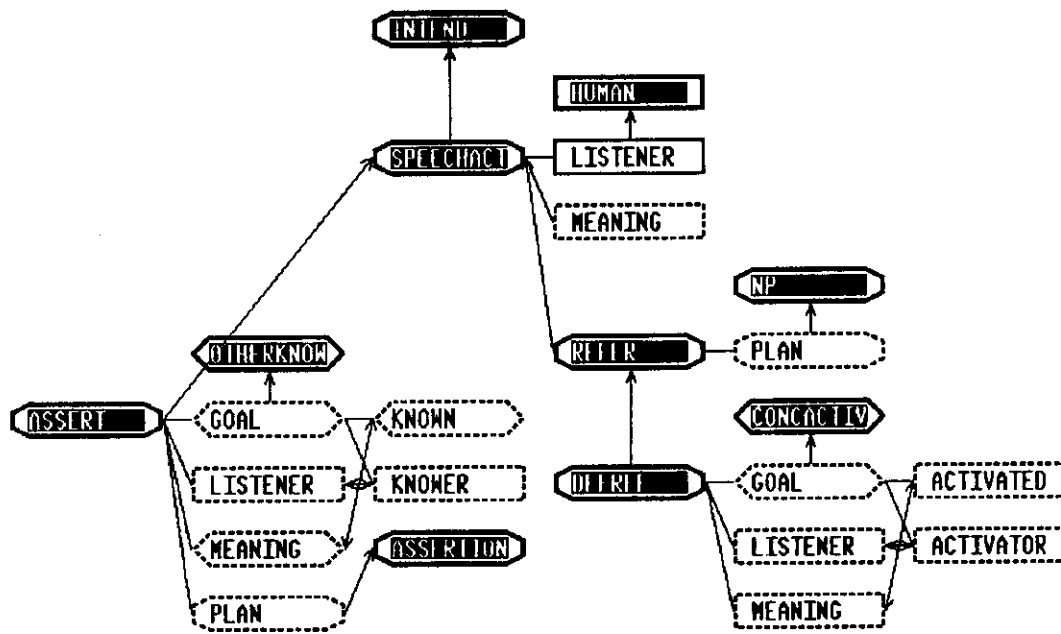


Figure 3

These correspond to goals 1. and 2. in the list in the last paragraph. Both are subtypes of the type SPEECH-ACT, from which they inherit a MEANING and a LISTENER role. ASSERT has a GOAL that another person KNOW a particular fact (the KNOWN role of the GOAL). This fact is also the MEANING of ASSERT. The PLAN for achieving the GOAL is a type of pattern called an ASSERTION. For DEFINITE-REFER the GOAL is that a particular concept be in an ACTIVE state in the mind of the LISTENER (see Appelt, 1985 for the notion of concept activation in reference). This concept becomes the MEANING, and the associated PLAN is an NP. Thus patterns are seen as PLANS for achieving communicative GOALS; this is why they are represented as ACTS rather than THINGS.

Figures 4, 5, 7 and 8 show how the classification mechanism works in the production of the utterance *this tea is strong* by a native speaker of English. The process begins (Figure 4) with an INTENTION (INTEND1) that MIKE, the listener, KNOW the fact SOLUTION-STRONG1.

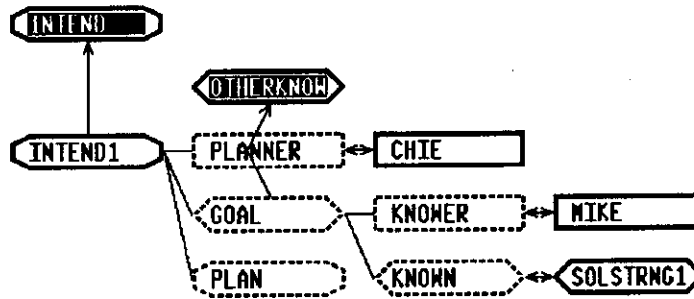


Figure 4

Initially the PLAN role of INTEND1 is not specified. Next INTEND1 is classified with respect to its GOAL role. Since this GOAL is an instance of the type OTHER-KNOW, it matches the GOAL role of ASSERT (Figure 3), and INTEND1 can be assigned to ASSERT. Now INTEND1 can inherit ASSERT's roles and the roles of any of ASSERT's ancestors that it is not already assigned to, in this case, SPEECH-ACT. Figure 5 shows the situation at this point.

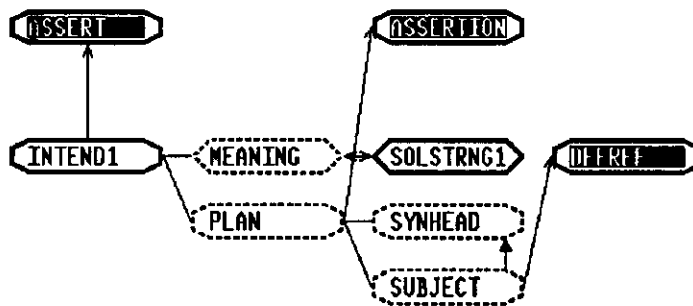


Figure 5

INTEND1 has inherited a MEANING role from SPEECH-ACT. This is equivalent to the fact that the speaker (PLANNER) wants the listener to know, SOLUTION-STRONG1. The PLAN role of INTEND1 has inherited from the PLAN roles of ASSERT and SPEECH-ACT the specification that the PLAN is an ASSERTION. From ASSERTION (and its ancestors), the PLAN has inherited a SUBJECT and a SYNTACTIC-HEAD role. The SUBJECT has already been assigned to DEFINITE-REFER; I will not go into how this classification is carried out. Note that the SUBJECT is not a pattern but rather a descendant of INTEND (and SPEECH-ACT) with its own GOAL and PLAN roles (not shown in Figure 5). It is the SUBJECT's PLAN role, not yet specified, which will be realized as a pattern.

The classification of the INTEND instance up to this point may be termed **pragmatic**; it is on the basis of the type of input GOAL alone. It provides some of the syntax of the emerging utterance (the PLAN of the INTEND instance) but of course does not yield an utterable PLAN. The instance must be further classified on the basis of the semantic content of the input, that is, on the basis of the MEANING role, which in this example is equivalent to the fact SOLUTION-STRONG1. Now the system searches for another descendant of SPEECH-ACT with a MEANING role which matches that of the input, in other words, for a suitable lexical entry. Each lexical entry in the system has a MEANING role which is at least partially specified and a PLAN role (for convenience, referred to as PHRASE) which is also partially specified. Figure 6 is a simplified view of a lexical entry, the one that is needed in the present example. The names of lexical entries begin with a colon.

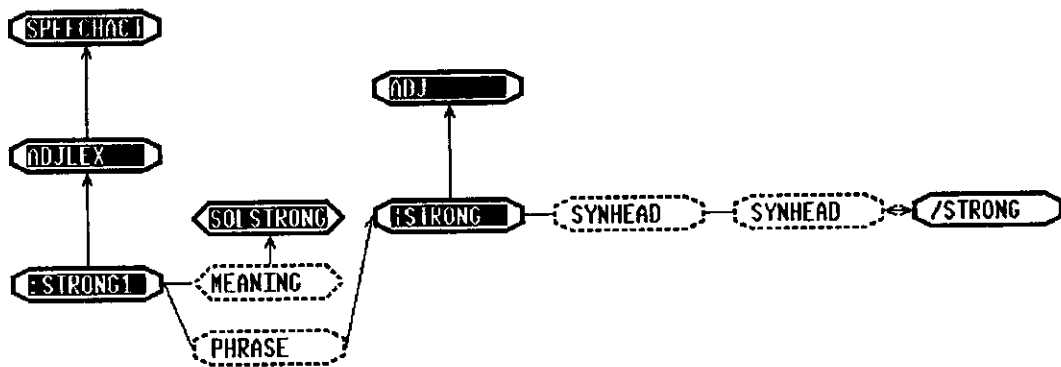


Figure 6

The significance of the entry is that there is a particular category of SPEECH-ACT which has a MEANING that is a fact of type SOLUTION-STRONG and a PHRASE (a PLAN to realize a communicative GOAL) that is an instance of |STRONG, an ADJECTIVAL pattern containing the morpheme /STRONG. Seen another way, this lexical entry says that a normal plan for talking about an instance of SOLUTION-STRONG is to use an instance of the pattern |STRONG. Because SOLUTION-STRONG1 is an instance of SOLUTION-STRONG, INTEND1 can be assigned to :STRONG1 and inherit from it. This allows a further specification of the PLAN role of INTEND1, as shown in Figure 7.

There are several more things that need to be done before the utterance is complete; they are outlined only briefly in what follows. First,* the instance of TEA which is the argument of SOLUTION-STRONG1 must be bound to the MEANING role of the SUBJECT of INTEND1. The information that this association is the appropriate one is a part of the lexical entry :STRONG1 but was omitted above for simplicity. Once the SUBJECT has a MEANING, it can be classified semantically. In this way, the lexical entry :TEA is found. In addition, the PLAN role of INTEND1 must be further classified. The types ASSERTION and ADJECTIVAL lexical pattern, both ancestors of this PLAN, have a common descendant ADJECTIVE-ASSERTION,

*For the sake of discussion, the steps in the formulation of the sentence are treated here in a particular sequence, but it is assumed that processing actually takes place in parallel (Bock, 1982; Dell, 1985; Kempen & Hoenkamp, 1982). This feature is not relevant to the concerns of this paper, however.




Figure 7

which has a role in its SYNTACTIC-HEAD for a COPULA. When the PLAN is assigned to ADJECTIVE-ASSERTION, the COPULA role can be inherited. Finally, on the basis of temporal properties of SOLUTION-STRONG1, INTEND1 can be assigned to PRESENT-ASSERT, which provides the appropriate tense for the COPULA. Figure 8 shows the utterance at its completion.

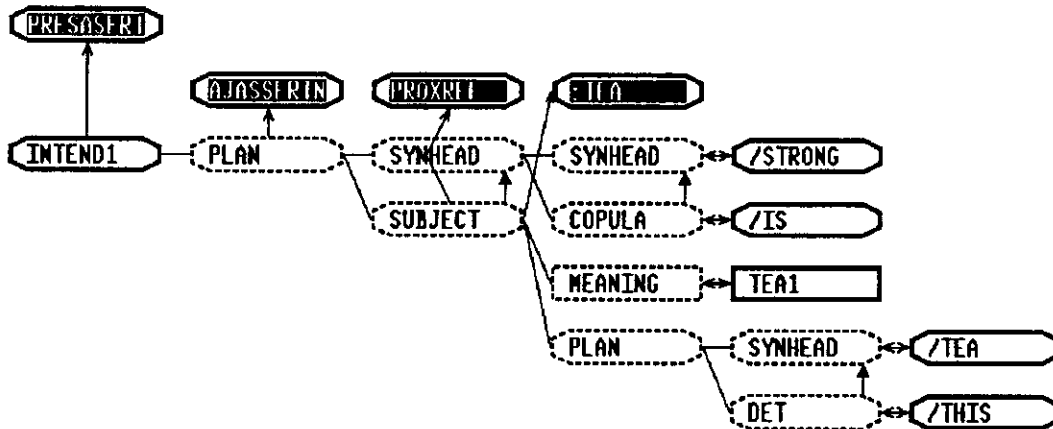


Figure 8

The generation of a sentence normally involves one or more subgoals in addition to the top-level goal. For instance, in ASSERTING a fact, as we have seen, there is usually a goal to REFER to one of the arguments of the fact, more precisely, a goal that the listener either find the referent in memory (*activate* it, in Appelt's, 1985 terms) or create a new node in memory for the referent if it is new. In the last example, there is a goal to have the speaker find the referent of *this tea* in memory in order to associate the new fact with something already known. The normal plan for such a goal is to produce a definite noun phrase for the referent. Another important type of subgoal arises when the speaker decides that there is the need to have the listener attribute a particular fact or property to a referent either because the attribute is interesting or because it is necessary for distinguishing the referent from other reference competitors (see Givón,

1983, for more on reference competition). The normal plan when such a situation arises is to generate a modifier of some sort, its precise form determined by the lexical pattern that is selected.

Not all pattern selection is made in response to explicit goals of the type described so far. Every language requires a speaker to make a number of decisions as side-effects of the selection of particular pattern types. For example, in English the selection of a lexical pattern belonging to the class COUNTABLE-NOMINAL forces the (native) speaker to make a choice between the SINGULAR and PLURAL patterns. This selection is normally unrelated to any explicit goal to convey the number of the referent, and of course in many languages, for example, Japanese, is not normally obligatory or even natural. The addition of another link type, the exhaustive-classification link, to the representation system outlined thus far will accommodate such decisions. This link is a variant of the is-a link; it joins a type concept to a set of its children which exhaustively classify it. The significance of the link is that every instance of the higher-level type must also be an instance of exactly one of the children. Each time a new instance of the higher-level type is created, there is automatically a check to determine which of the children it should be assigned to. One of the children may be designated the default; if the new instance matches none of the other children, it must be assigned to the default. Figure 9 shows how English noun number is represented. The exhaustive-classification link, indicated by the small hexagon with the X in it, connects :COUNT with two of its children, :SINGULAR and :PLURAL. Whenever an instance of :COUNT is created, for example, when referring to an instance of the concept BASKET and matching the lexical entry :BASKET, the classification mechanism would on the basis of the MEANING role then assign the instance to either :PLURAL or :SINGULAR. :PLURAL is selected if the referent is a COLLECTION. Otherwise :SINGULAR, the default, is selected.*

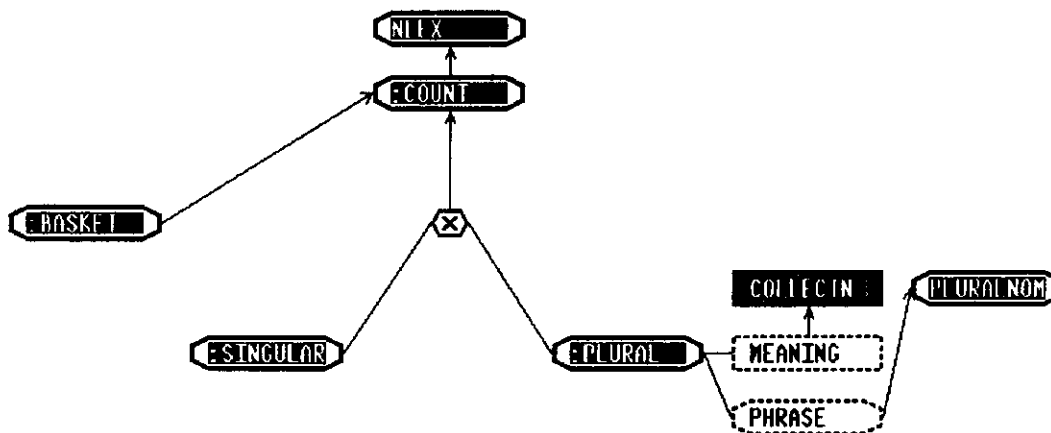


Figure 9

*Exhaustive classification does not suffice to yield all grammatical morphology. Grammatical gender agreement, for example, requires a further notion, that of the greatest common descendant of two or more types. As with exhaustive classification, however, there is no reliance on the input goal for the selections that are made.

In summary, language production is viewed as the selection and combination of linguistic patterns primarily on the basis of the pragmatic and semantic content of input goals and to a lesser extent on the basis of distinctions made by the language that are unrelated to input goals. The discussion here has covered only those issues in the computational approach to language production that are pertinent to the types of errors being considered. For other treatments, the reader is referred to Appelt (1985), Kempen & Hoenkamp (1982), Mann & Matthiessen (1983), McDonald & Pustejovsky (1985), and McKeown (1985).

4. Gaps in Linguistic Knowledge

Consider now at what points in the system described in 3.3. and 3.4. a language learner might be deficient. The speaker who produced (1) lacks* knowledge that a count noun in English must be plural if its referent is a set of things and knowledge that a verb in a present clause takes a suffix if its subject is third person singular. In terms of the model, these problems result from the lack of an exhaustive classification under a high-level goal-pattern association. Thus for this speaker the lexical entry :BASKET (Figure 9) would be a child of the type NOMINAL-LEXICAL-ENTRY, and :COUNT and its subtypes would not yet exist. Because the missing associations such as these do not serve to satisfy any communicative goals, this sort of gap does not interfere with processing. In fact, the speaker is not normally even aware that such a gap exists.

In (2), (3), and (4), on the other hand, communicative goals are involved in the errors that are made, and thus it makes sense to deal with these in terms of planning. In (2), repeated here for convenience, the speaker is obviously experiencing difficulty finding a lexical item for the event in memory, and she hesitates to proceed until she is sure that the listener understands her.

- (2) And uh passing away each other pass away
 no uh passing passed
 [Native speaker: Left each other?]
 Left each other.

The event that the speaker intends to convey, call it EVENT1, could be represented as an instance of the concept SEPARATE, or perhaps something more specific such as LEAVE-ONE-ANOTHER-AFTER-MEETING. The problem, in terms of the model, is that there is apparently no English lexical entry for these type concepts, and the speaker must somehow arrive at a lexical pattern which can do the job. While I will not consider this example in detail, the relevant strategy in such cases appears to be to search through higher levels of the conceptual hierarchy to find a concept that has a lexical entry. Thus one ancestor of EVENT1 is the concept RECIPROCAL-ACT, for which the speaker apparently has a lexical entry associating it with the lexical pattern |EACH-OTHER.

In (3), repeated here, there is no way of telling whether we are dealing with a faulty association that has already been learned or with a selection that is made during production in response to a gap in knowledge.

- (3) I don't like this tea because it's too thick (= strong).
 (Takahashi, 1985)

*CHIE does not yet allow for partial knowledge, that is, knowledge that is in memory, but relatively weak and inaccessible. Partial knowledge would help to account for apparent free variation. One way to implement this notion would be to add the concept of node and/or link strength to the model (as in Anderson, 1983, for example).

If the latter is the case (and even if it is not, comparable situations certainly do arise in second language production contexts), the problem is again the lack of a lexical entry, here for the concept SOLUTION-STRONG. Let us consider this example in more detail, restricting ourselves to the production of the simpler sentence *this tea is thick*. Up to a point, the formulation process is identical to what we saw in the example of a native speaker's production of the corresponding correct sentence, except that now classification proceeds by language as well as by MEANING and GOAL type. Thus it is assumed that each instance of SPEECH-ACT and each pattern has a LANGUAGE role. Figure 10 shows the state of affairs for the present example corresponding to that in Figure 5. ENGLISH-ASSERT, ENGLISH-ASSERTION, and ENGLISH-DEFINITE-REFER are children of ASSERT, ASSERTION, and DEFINITE-REFER respectively.

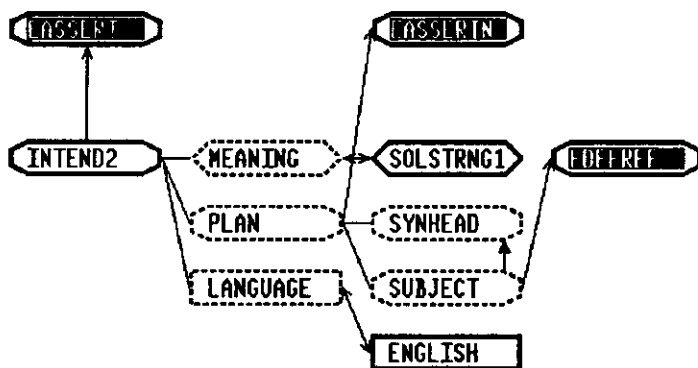


Figure 10

The difficulty arises when the speaker must classify the speech act by its MEANING. The concept which is equivalent to the MEANING, SOLUTION-STRONG1, is represented as an instance of the fact type SOLUTION-STRONG, which in turn is a child of the more general type CONCENTRATED. CONCENTRATED has another child VISCOUS. The distinction between the two subtypes of CONCENTRATED is in the way in which the concentration of the solution manifests itself, for VISCOUS in terms of the resistance of the solution to being stirred or poured quickly and for SOLUTION-STRONG in terms of the effect the solution has upon other objects. CHIE has a Japanese lexical entry for CONCENTRATED with the PHRASE [KO- and an English lexical entry for VISCOUS with the PHRASE [THICK, but no English lexical entry for SOLUTION-STRONG. Figure 11 illustrates this portion of the lexicon. In classifying INTEND2 semantically, CHIE looks at both the MEANING and LANGUAGE roles of potential matches. Two *partial* matches are found, J:CONCENTRATED and E:VISCOUS. For J:CONCENTRATED there is a match on the MEANING role because the fact which is equivalent to INTEND2's MEANING, SOLUTION-STRONG1, is a descendant of CONCENTRATED. However, the LANGUAGE role of J:CONCENTRATED, which is equivalent to JAPANESE, fails to match the LANGUAGE role of INTEND2, which is equivalent to ENGLISH. Conversely, for E:VISCOUS there is a match on the LANGUAGE role but the MEANING role fails to match because a solution which is VISCOUS is not necessarily SOLUTION-STRONG.

Given the two partial matches, one strategy is to create a new tentative lexical entry that combines the matching features. The entry which CHIE creates for this example is shown in Figure 12. The new entry, E:CONCENTRATED, has the MEANING role of J:CONCENTRATED and the LANGUAGE and PHRASE roles of E:VISCOUS. This tentative entry provides a match for INTEND2, and the formulation of the sentence can

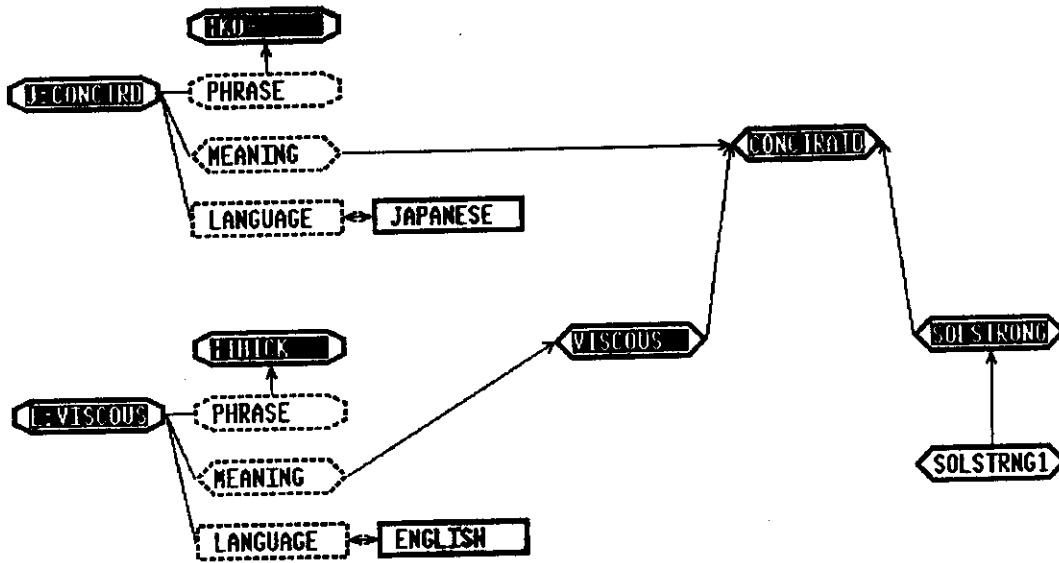


Figure 11

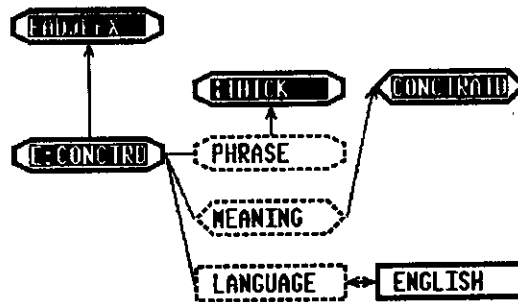


Figure 12

continue as in the example in 3.4., except that we get /THICK in place of /STRONG. The strategy that applies in this situation is the following:

Given: a SPEECH-ACT instance *s*
 with MEANING *M*
 and LANGUAGE *L2*

If there is no lexical entry which matches both *M* and *L2*,
 and
 there is a lexical entry *LE1* with MEANING *LEM1*

matching M,
 and
 there is a lexical entry LE2 with LANGUAGE = L2 and
 MEANING LEM2 a descendant of LEM1

 then create a tentative lexical entry with
 MEANING = LEM1
 LANGUAGE = L2 and
 PHRASE = PHRASE of LE2,
 and
 assign s to this entry (inheriting PHRASE of LE2)

Note the condition that the MEANING of the partially matching L2 lexical entry be a descendant of the MEANING of the partially matching L1 lexical entry. In other words, in addition to the match on the LANGUAGE role of the L2 entry, there must be at least a partial match on the MEANING as well.

In (4), repeated here with irrelevant errors corrected in the English version, the problem appears to be the unavailability of the English relative clause pattern to the speaker.

(4a) Soshite sono teono o otoshita otoko ga
 then that hatchet OBJ dropped man SUBJ
 'Then the man who dropped the hatchet

sono teono o sagashite imashita.
 that hatchet OBJ was seeking
 was looking for the hatchet.'

(4b) And the hatchet was somebody's.
 (Somebody had dropped the hatchet.)
 He was looking for his hatchet.

In this context the need to produce a relative clause arises because of the presence of three men in the story. Without the attribution of an identifying fact to the referent of the noun *man*, the listener will have difficulty identifying the referent.*

As one of the children of the pattern NOMINAL, CHIE has a pattern MODIFIED-NOMINAL with a MODIFIER role. A MODIFIER is an instance of the type ATTRIBUTE, a SPEECH-ACT subtype distinct from the two we have already encountered, ASSERT and REFER. MODIFIED-NOMINAL has three subtypes which are distinguished by what type of lexical entry the MODIFIER is assigned to. ADJECTIVE-NOMINAL has a MODIFIER which is an ADJECTIVE-LEXICAL-ENTRY, POSSESSIVE-NOMINAL has a MODIFIER which is a NOMINAL-LEXICAL-ENTRY, and RELATIVE-CLAUSE-NOMINAL has a MODIFIER which is a VERB-LEXICAL-ENTRY. Figure 13 illustrates how some of these concepts are related. In producing a noun phrase such as the subject of (4a) the PLAN role of the REFER instance is first assigned to MODIFIED-NOMINAL and then further classified according to what sort of lexical entry is selected for the MODIFIER. Because the three children of MODIFIED-NOMINAL are realized differently in English and Japanese, they are further classified by LANGUAGE. CHIE has Japanese subtypes for all three but is lacking an English subtype for RELATIVE-CLAUSE-NOMINAL. The left side of Figure 14 shows some of these subtypes. Note that the JAPANESE-

*The planning necessary to assess this difficulty is quite complex and has not yet been worked out in detail.

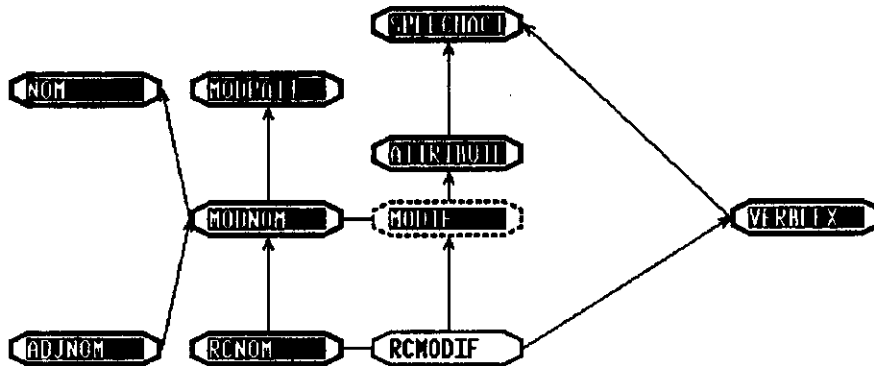


Figure 13

RELATIVE-CLAUSE-NOMINAL description specifies that the RELATIVE-CLAUSE precede the SYNTACTIC-HEAD of the pattern.

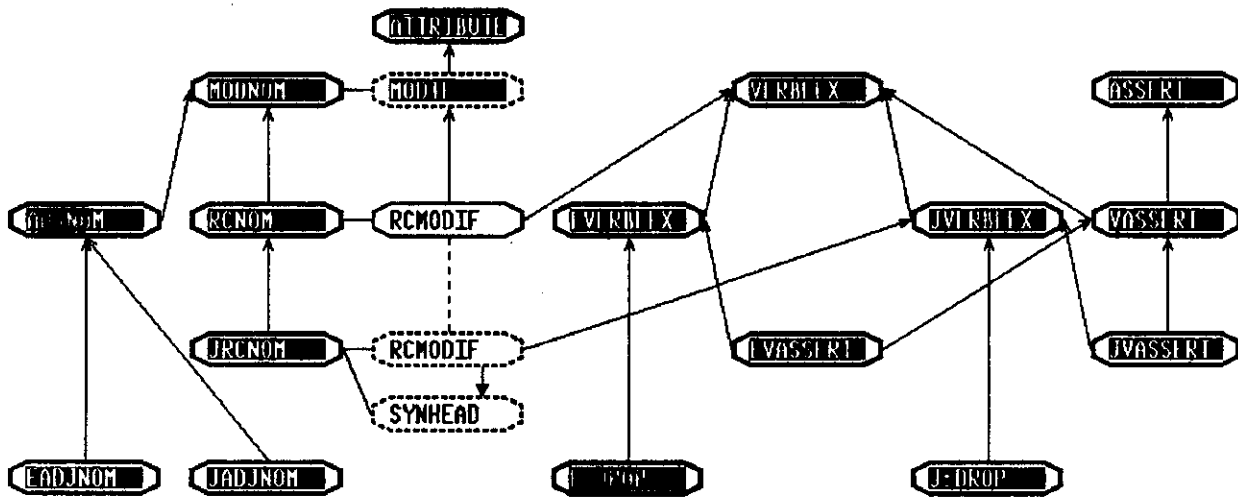


Figure 14

The right side of Figure 14 shows how relative clauses are integrated into the lexicon. A Japanese lexical entry such as J:DROP, the entry which is selected for the relative clause in the subject of (4a), is a child of JAPANESE-VERB-LEXICAL-ENTRY, which has three possible syntactic/pragmatic realizations, as a JAPANESE-VERB-ASSERT, a JAPANESE-VERB-REFER (not shown in the figure), or a RELATIVE-CLAUSE in a JAPANESE-RELATIVE-CLAUSE-NOMINAL. An instance of J:DROP, then, can serve to either ASSERT or ATTRIBUTE the fact that it refers to. The corresponding English lexical entry E:DROP, on the other hand, is more limited. Its parent, ENGLISH-VERB-LEXICAL-ENTRY, has only one syntactic/pragmatic subtype, ENGLISH-VERB-ASSERT. Thus an instance of E:DROP can only be used to ASSERT a fact.

Consider how this works in the production of (4a) and (4b). Let us assume that the speaker has chosen the fact referred to by the sentence *the man dropped the hatchet* to attribute to the referent of the noun phrase being formulated. I shall call this fact DROP1. Figure 15 shows one stage in the production of the noun phrase in question in Japanese.

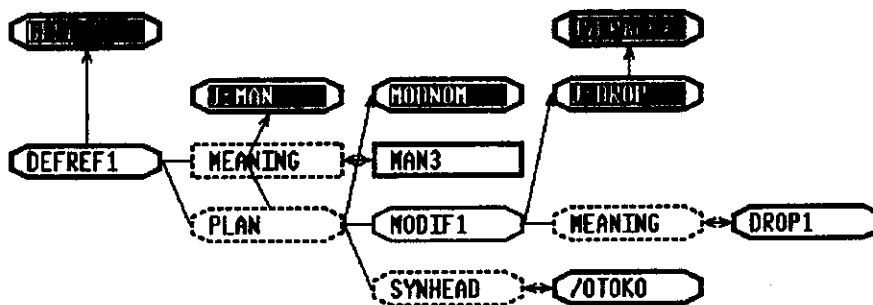


Figure 15

The REFER instance has already been classified as DEFINITE, and the lexical J:MAN has yielded the SYNTACTIC-HEAD /OTOKO. A MODIFIER role has been created, and it has been assigned to the lexical entry J:DROP on the basis of its MEANING. Now the PLAN of DEFINITE-REFER1 must be further classified so as to specify what precise form the MODIFIER will take and how it will be ordered with respect to the SYNTACTIC-HEAD. The PLAN matches JAPANESE-RELATIVE-CLAUSE-NOMINAL because its MODIFIER is a descendant of JAPANESE-VERB-LEXICAL-ENTRY. In Figure 16 the PLAN has been assigned to JAPANESE-RELATIVE-CLAUSE-NOMINAL, and the constituent ordering has been inherited.

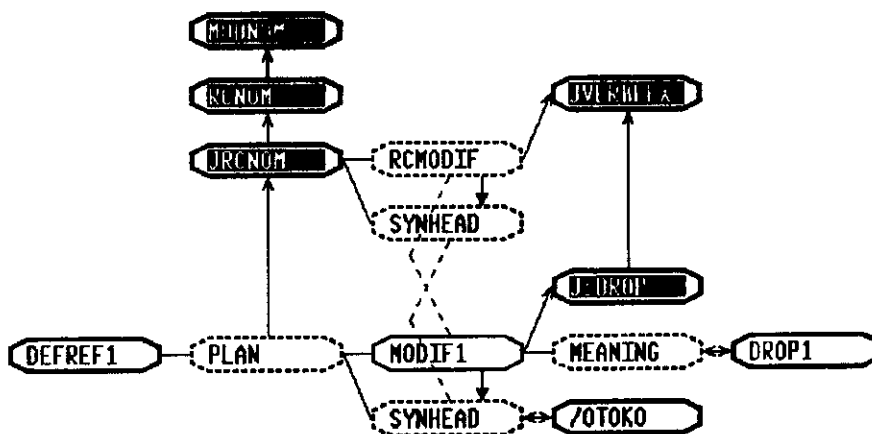


Figure 16

In attempting to produce the corresponding noun phrase in English, CHIE begins in the same way. Figure 17 illustrates the stage in the process corresponding to that in Figure 15. Again there is the problem of what form the MODIFIER is to take and how it is to be integrated into the noun phrase. Now, however, there is no subtype of MODIFIED-NOMINAL to do the job.




Figure 17

CHIE is left with two SPEECH-ACTS, DEFINITE-REFER2 and MODIFIER2, neither of which is utterable in its present form. In such a situation there are several possible strategies that a speaker might resort to. One is to abandon one of the GOALS embodied in the two SPEECH-ACTS. Since MODIFIER2 has arisen as a result of a subgoal generated within DEFINITE-REFER2, it is this SPEECH-ACT that would be sacrificed. This would result in a sentence of the form *the man was looking for the hatchet* with no fact attributed to the man in question. The decision to abandon the attribution goal, like the original decision to make the attribution, is a complicated one, and the model in its present form does not offer this option. It would be interesting to determine whether such potentially ambiguous references are more frequent in the speech of second language learners than in that of native speakers. Another possible strategy is to attempt to achieve the attribution goal by using as much of the available L1 pattern as possible; in this context this would mean placing the relative clause before the head of the noun phrase and omitting the relative pronoun which English requires. Though this does occur in the speech of Japanese learners of English, it is again not what we see in (4b).

The strategy proposed here to account for this error is to attempt to accomplish at least partially both SPEECH-ACTS. This is done by assigning MODIFIER2 to the only available descendant of SPEECH-ACT which specifies a pragmatic/syntactic pattern, ENGLISH-VERB-ASSERT. The situation at this point is shown in Figure 18. The PLAN of MODIFIER2 is then realized as the first sentence in (4b). Now the PLAN of DEFINITE-REFER2 can no longer be realized as a MODIFIED-NOMINAL, and it takes the form of a pronoun or an unmodified noun preceded by a definite article. Note how the result deviates from what the speaker intended. The assertion in the first sentence is pragmatically inappropriate because it implies that the fact is new information though its semantic content is the same as that of the original GOAL. However, this makes the definite reference in the second sentence pragmatically appropriate: no modifier is required following the reference in the first sentence to the same man.

The strategy can be formulated roughly as follows:

Given: a SPEECH-ACT instance *s*

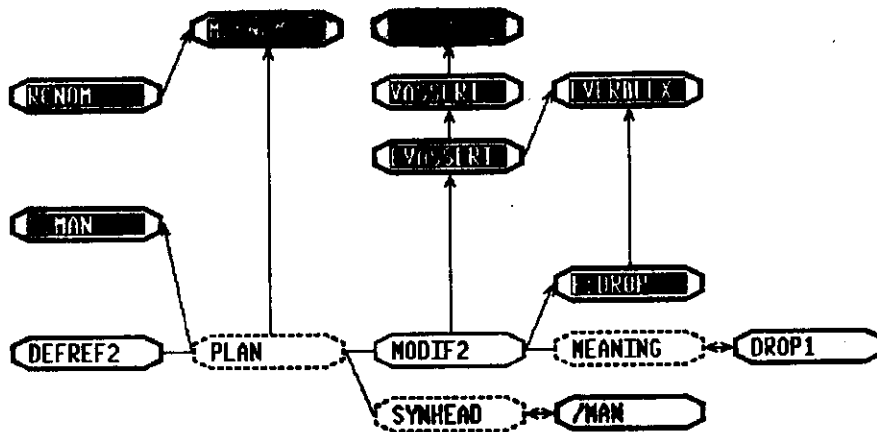


Figure 18

If there is an L2 lexical entry matching *s*
 and
 there is no L2 SPEECH-ACT descendant matching
 the GOAL type of *s*
 then assign *s* to another (pragmatically inappropriate)
 L2 SPEECH-ACT descendant

It remains to be seen how general this strategy is. For example, we might expect it to apply when a learner is not familiar with certain types of nominalizations. Thus in place of something like

The fact that the tea is too strong doesn't surprise me.

we might expect to find

The tea is too strong.
 This doesn't surprise me.

5. Future Work and Conclusions

The work reported on of course only lays the groundwork for the development of a full-scale theory of second language production. The model is currently being extended to

1. include a wider range of strategies, in particular, the possibility of goal abandonment,
2. provide for the selection of one strategy when several are available, as with the relative clause problem in (4),

3. better deal with the decisions involved in reference,
4. characterize snapshots of several points in the acquisition process.

In addition, an attempt is being made to relate the production strategies that have been isolated to

1. L2 comprehension strategies,
2. generalization strategies recognized in cognitive psychology and machine learning research,
3. processes occurring in creolization.

In this paper I have argued that a characterization of the mental representation of linguistic knowledge and of the cognitive processes involved in language production is crucial to a theory of second language acquisition. I have shown how several examples of errors made by language learners can be accommodated in a model of language production as the realization of communicative goals through the selection and combination of linguistic patterns from a network of patterns. Learners err when there is a gap in their knowledge of the target language, but only when this gap is related directly to the realization of a communicative goal are they forced to deal with it.

References

- Andersen, R. W. (ed.). 1983. *Pidginization and Creolization as Language Acquisition*. Rowley, Mass.: Newbury House.
- Anderson, J. R. 1983. *The Architecture of Cognition*. Cambridge, Mass.: Harvard University Press.
- Appelt, D. E. 1985. *Planning English Sentences*. Cambridge: Cambridge University Press.
- Becker, J. D. 1975. The phrasal lexicon. In *Proceedings of the Interdisciplinary Workshop on Theoretical Issues in Natural Language Processing*. Cambridge, Mass., June 1975.
- Brachman, R. J., and Schmolze, J. G. 1985. An overview of the KL-ONE knowledge representation system. *Cognitive Science* 9, 171-216.
- Bresnan, J. W., and Kaplan, R. M. 1983. Lexical-functional grammar; a formal system for grammatical representation. In J. W. Bresnan (ed.), *The Mental Representation of Grammatical Relations*. Cambridge, Mass.: MIT Press.
- Brumfit, C. J., and Johnson, K.. 1979. *The Communicative Approach to Language Teaching*. Oxford: Oxford University Press.
- Dell, G. S. 1985. Positive feedback in hierarchical connectionist models: applications to language production. *Cognitive Science* 9, 3-23.

- Fahlman, S. E. 1979. *NETL: A System for Representing and Using Real-World Knowledge*. Cambridge, Mass.: MIT Press.
- Fillmore, C. J. 1979. Innocence: a second idealization for linguistics. *Proc. 5th Ann. Meeting of the Berkeley Ling. Soc.*, 63-76.
- Flickinger, D., Pollard, C., and Wasow, T. 1985. Structure-sharing in lexical representation. *Proc. 23rd Ann. Meeting of the Association for Computational Linguistics*, 262-267.
- Fodor, J. A. 1983. *The Modularity of Mind*. Cambridge, Mass.: MIT Press.
- Gazdar, G., Klein, E., Pullum, G., and Sag, I. 1985. *Generalized Phrase Structure Grammar*. Cambridge, Mass.: Harvard University Press.
- Givón, T. (ed.). 1983. *Topic Continuity in Discourse: a Quantitative Cross-Language Study*. Amsterdam: John Benjamins.
- Hatch, E. M. 1983. *Psycholinguistics: A Second Language Perspective*. Rowley, Mass.: Newbury House.
- Hatch, E. M., Flashner, V., and Hunt, L. In press. The experience model and language teaching. In R. Day (ed.), *Conversation in Second Language Acquisition*, Rowley, Mass.: Newbury House.
- Herskovits, A. 1985. Semantics and pragmatics of locative expressions. *Cognitive Science* 9, 341-378.
- Jacobs, P. 1985. PHRED: A generator for natural language interfaces. University of California at Berkeley, Computer Science Division Memorandum UCB/CSD 85/198.
- Kay, M. 1979. Functional grammar. *Proc. 5th Ann. Meeting of the Berkeley Ling. Soc.*, 142-158.
- Kempen, G., and Hoenkamp, E. 1982. An incremental procedural grammar for sentence formulation. Unpublished paper, Department of Psychology, University of Nijmegen.
- Krashen, S. 1981. *Second Language Acquisition and Second Language Learning*. London: Pergamon Press.
- Mann, W. C., and Matthiessen, C. M. I. M. 1983. *Nigel: a Systemic Grammar for Text Generation*. RR-83-105, USC/Information Sciences Institute, February, 1983.
- McDonald, D., and Pustejovsky, J. 1985. Description-directed natural language generation. *Proc. 9th International Joint Conference on Artificial Intelligence*, Los Angeles, 799-805.
- McKeown, K. R. 1985. *Text Generation*. Cambridge: Cambridge University Press.
- McShane, J., and Dockrell, J. 1983. Lexical and grammatical development. In B. Butterworth (ed.), *Language Production, Volume 2: Development, Writing and Other Language Processes*. London: Academic Press.

- Nattinger, J. R. 1984. Communicative language teaching: a new metaphor. *TESOL Quarterly* 18, 391-407.
- Pawley, A., and Syder, F. H. 1983. Two puzzles for linguistic theory: nativelike selection and nativelike fluency. In J. C. Richards and R. W. Schmidt (eds.), *Language and Communication*. London: Longman.
- Pinker, S. 1984. *Language Learnability and Language Development*. Cambridge, Mass.: MIT Press.
- Schank, R. C., and R. Abelson. 1977. *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum.
- Selinker, L. 1972. Interlanguage. *International Review of Applied Linguistics* 10, 209-231.
- Stemberger, J. P. 1982. Syntactic errors in speech. *Journal of Psycholinguistic Research* 11, 313-345.
- Takahashi, T. 1985. Language transfer and the acquisition of patterns of lexico-semantic correspondence. Paper presented at the Sixth Los Angeles Second Language Research Forum, Feb. 23, 1985, Los Angeles, CA.
- Wilensky, R. 1983. *Planning and Understanding: a Computational Approach to Human Reasoning*. Reading, Mass.: Addison-Wesley.
- Wilensky, R., and Y. Arens. 1980. PHRAN--A Knowledge-based Approach to Natural Language Analysis. University of California at Berkeley, Electronics Research Laboratory Memorandum UCB/ERL M80/34.

