A HOLISTIC APPROACH TO LANGUAGE

Brian D. Josephson[1] and David G. Blair[2]

Cavendish Laboratory, Madingley Road,
Cambridge CB3 0HE, England.

Progress report no. TCM/32/1982
November 29th, 1982
(with corrections 1995)

Abstract

The following progress report, produced originally in 1982, views language acquisition as primarily the attempt to create processes that connect together in a fruitful way linguistic input and other activity. The representations made of linguistic input are those that are optimally effective in mediating such interconnections. An effective Language Acquisition Device should contain mechanisms specific to the task of creating the desired interconnection processes in the linguistic environment in which the language learner finds himself or herself. Analysis of this requirement gives clear indications as to what these mechanisms may be.

[1] *email*: bdj10@cam.ac.uk

[2] Present address: Maritime Operations Division, DSTO, PO Box 44, Pyrmont, NSW 2009, Australia. *email*: david.blair@dsto.defence.gov.au

1 Introduction

This progress report gives in a concise form the results of a new approach to the problem of language, one which attempts to deal with the subject matter on a holistic basis, that is to say we require that our ideas be applicable to the phenomenon of language as a whole, and not to deal with, say, just the linguistic aspects or just the psychological aspects. The latter type of approach is, we suspect, similar to trying to understand a three-dimensional object on the basis of a two-dimensional projection of it: theories obtained by studying a single aspect of language may work well within that particular domain but be quite inapplicable outside it.

We were fortunate in having available at the outset a little known general account of many aspects of intelligence, including language (Mahesh Yogi 1974). Some clarifications of this very abstract description were gained from the work of Skemp (1979) and that of Eckblad (1981). The need to integrate very general principles with attention to matters of detail was met by utilising ideas from the work of Dik (1978) and Marcus (1980).

Little can be done in this introduction to summarise the theory, except to say that the rules for processing language, and the mechanisms for their acquisition, play a central role; and further, that by keeping attention simultaneously on the levels of very abstract principles and of less abstract general facts about language one is led upon a subtle trail, on which one finds that all the general rules get explained in a essentially simple and direct manner.

The general form of this work involves types of argument very different from the usual analytical-deductive ones (which contrive to 'prove' a given hypothesis to be the correct one), and many of the ideas are contained implicitly in the arguments rather than being stated explicitly. Exclusively left-brain thinkers should perhaps stop reading right away. For the remaining readers, it is suggested that more attention should be given to absorbing the basic ideas than to their critical analysis. It is important that the mind be allowed to work on the ideas spontaneously. For ourselves, we find the degree of fit that can be obtained between the generally intuitively appealing hypotheses and the patterns in the linguistic and other data to be a good sign that the ideas are a reasonable approximation to the truth. It is not feasible at this stage in the research to attempt to make the theory perfect and to study exhaustively all permutations and combinations of the concepts involved; therefore there are an number of gaps in the theory as formulated in this report, and it should therefore be viewed as a convenient stopping place on the way to perfection (which is a very long way away from where we are now). The present work can be only a first approximation to the truth, and future research will be needed to amend any errors and add further necessary detail.

2 Language in General

A linguistic message is exchanged between sender and receiver, via a communication channel. We shall concentrate here on the processes of the receiver, for two reasons: (i) understanding language does not depend much on the ability to speak, and so it is legitimate to consider it in isolation and (ii) understanding is primary, in the sense that the speaker has to adapt what he says to what the listener is able to understand, but the converse does not apply. (Our concentration on the receiver is due only to our wish to limit the breadth of the problems under consideration, and not to any fundamental inapplicability of our concepts to the problem of production of speech).

Linguistic exchange is a process of great utility to the users of a language. This utility is contingent upon the receiver being able to translate the message into a format in which it is able to guide his actions. It must be translated from the code of the language into an inner code before this is possible. This we refer to as the process of interpretation. The translation process and processes auxiliary to it involve many specific subprocesses, which are learnt during the acquisition of the language concerned. Language learners use only some of these processes, and they may use simplified or erroneous processes which lead to failure to understand or misunderstanding.

The inner code, in order that it can guide actions, must in some sense consist of the individual's representation of what is or what may be (or what was). The existence of inner code is independent of language; language is merely an alternative way of generating it. Language acquisition is in essence learning the right way to connect inner code and external language. But there is no unique 'right way' since in reality language users have a diverse variety of processing systems. Their environments differ and so what they acquire is different. But their environments also have similarities, and different users' systems are in general equivalent enough for effective information transfer to take place. Nevertheless, the 'meaning' assigned to an utterance is in reality (except in special cases) a subjective entity (and to a degree an intersubjective one), rather than an objective one as assumed in many treatments of language.

Grammar plays an important but not exclusive role in regulating the initiation of the interpretative processes. Inappropriate input is ignored in the first instance, greatly reducing the number of possibilities that are considered by the listener. The appropriateness of input is governed mainly by linguistic categories. By selecting the categories in his linguistic output, as a speaker of a language learns to do, a speaker can determine which combinations of entities a listener will assemble.

Language has a pragmatic as well as a semantic component (as well as syntactic and phonological components). That is to say, certain linguistic forms presuppose that the listener will do certain things with the information supplied. Users learn the advantages of doing these particular things, embodying the information in their system of rules or processes, while the fact that listeners know and apply these conventions is used by speakers when they make use of language.

3 The acquisition and application of processes (or rules)

The treatment given here resembles in many ways the theory of Skemp (1979). Note first that a process which a language user applies in processing language is effectively the same as a rule telling him what he should do. This kind of rule is not the same as a grammarian or linguist's rule, the latter rules being *descriptive* while the rules relating to actual processes are *prescriptive*. But the two kinds of rules are clearly very closely linked. For example, a linguist's rule that a certain word (in a given context, perhaps) has a certain meaning carries with it a prescriptive rule that if the word is heard in the given context then a code which represents that meaning should be created by the system. From this we conclude that the core of the acquisition system should be a set of processes which take note of relationships of the type known to the linguist, and automatically modify the system so that the corresponding process becomes available (or rule is applied).

Having noted this, we must go on to observe that any rules discovered by such means may be wrong, or valid only under certain conditions. Therefore the system must go on to learn (and apply) further knowledge which restricts the application of the rules. The system must learn two distinct things, which we shall term the *domain* of applicability of the rules and their *conditions* of applicability. The sense of the former is that of the domain of a function in mathematics, in other words it is only within a certain set of entities that the rule should be applied at all, while the question of whether a rule should be applied to an entity which is inside the domain is a matter of other conditions which may apply independently of the entity being operated upon by the rules. From the point of view of learning the domains and conditions it is crucial to observe that when a rule is applied to an entity outside the domain of the rule, a significant outcome can result only by chance, while a greater than chance significant outcome will occur for an entity within the domain (depending on the probability that the appropriate *condition* is satisfied). In principle, then (we do not wish to consider details here), a system can be designed to discover the domain and conditions of a rule. Note that if an erroneous rule is by chance acquired, the criteria given will lead to the domain becoming non-existent.

For completeness, it should be noted that a system may extend its concept of the domain of a rule to take into account the dependence of a rule on a parameter, that is to say that the domain is parameter dependent. The system extends its first approximation of a fixed domain to take into account this parameter dependence. This point is exemplified by the fact that the rules for assigning a phoneme to a sound are context dependent and also dependent on factors such as the accent of the speaker. The mechanics involved will not be discussed further here.

Our approach demands the existence of processes which produce modifications or variants of rules. A number of such rule-changing or variant creating processes are (i) making a new rule in which a constant in the old rule becomes a variable (whose domain of variation is subsequently determined); this is a form of generalisation; (ii) adding extra variables: this applies when the system learns a construction slightly more complex than one it can already handle; (iii) making a mutant rule with some constant changed; (iv) rules made by 'process splitting', a mechanism which involves partitioning a previously learnt process into two components. Examples will be given later.

The mechanisms just discussed can give rise to the kind of structure noted in systemic grammar. The mutant process is one mechanism for this. Also it may be the case that one of the conditions for one rule is that another rule should be operating at the same time; this gives rise to the embedding of one process in another.

4 Language acquisition as evolution

If the system continually adds and elaborates its rules, and at the same time it operates a preference for those rules which produce the best outcome (when applied to the understanding of a sentence and making use of the information acquired), it will gradually approach the optimum rule structure, that of the expert user of the language. We note a regulatory mechanism which applies here: if a structure is difficult to learn, speakers will tend not to use it, since they will have difficulty being understood. But speakers do invent new processes (words, applications of words, occasionally new grammatical constructions). Whether these are propagated into the language or not is a function of their usefulness, through fairly obvious mechanisms which will not be detailed here.

5 The Functional Approach of Dik and its interpretation within the current framework

Dik's approach to language is to regard the verbal input as supplying information needed by the listener, according to a set of well defined conventions. For example, if the listener is being required to move something, he needs to know (a) the nature of the task, indicated by the word 'move' (b) the values of certain relevant variables, such as the object, its destination and the means employed. In this example the verbal input has a general form such as 'move X to Y, using Z', where X, Y and Z denote linguistic expressions. How does such an utterance create its effects? We seek a simple mechanism consistent with the facts.

We consider first the effect of the word 'move'. Our model is that (at a certain level of its functioning, at any rate) the system first creates a code $\lambda(move)$ corresponding to the word move. It subsequently translates this code into a code $\sigma(move)$ representing the *meaning* of move, i.e. the process of moving something. The second, semantic code, is one which under the right conditions is capable of activating programs for moving objects. Such programs will require at particular points information such as indicated by the linguistic expressions X, Y and Z. When the utterance is heard and interpreted, this information must be stored away (in either long-term or short-term memory) in such a way that it automatically reappears when it is needed.

The memory problem is one not specifically connected with linguistic input. The system has just as much the ability to remember actions which it has discovered for itself by trial and error. We assume that collections of information are stored in independently addressable memory blocks, and that within a block specific items are stored and retrieved by means of specific 'enquiry codes'. If this is the case, it is natural to assume that information acquired linguistically is stored in memory blocks using these same enquiry codes.

But what codes are stored exactly corresponding to X, Y and Z in our example? Requiring that the elementary constituents of processes be as simple as possible leads us to the hypothesis that the system uses a set of codes x, y and z which act as codes not only for the *linguistic form* of the unit (X, Y or Z) but also for the *meaning*. The codes x, y and z act as labels for blocks of memory which contain all the relevant information. The assumption is that when a linguistic unit such as that represented by X is recognised a label is assigned to it, with a corresponding block of memory, which is used to store, with the aid of suitable enquiry codes, the information indicated. Interpreting units such as X, Y and Z is a matter of taking information out of a given block of memory, transforming it in accordance with the relevant interpretation rule, and inserting the result back in the same block of memory (but not in identical memory locations since the enquiry codes for the different types of information are different).

The nature of a typical interpretation rule or process may be seen by returning to our example 'move X to Y, using Z'. The input to the process is the collection of codes $\lambda(move)$, x, y, z (stored under suitable enquiry codes), and the output is the collection $\sigma(move)$, x, y, z, again stored under suitable enquiry codes. The operation involved is the translation of the code for *move* from the λ version to the σ version, together with the copying of the codes x, y and z from their original enquiry code locations into new enquiry code locations associated with semantic information.

The effect of the interpretation process is to produce a set of information codes, in turn marked by enquiry codes, held either in short term memory or in a labelled block of long-term memory. The system must be such that it can treat the stored codes such as x, y and z as labels for subroutines and make the necessary arrangements for returning at the end. The most plausible mechanism is one which conforms to the intuitive idea of making a specific plan of action out of available knowledge (cf. Eckblad's discussion relating to the computer language SIMULA 67 (p. 104 of Eckblad 1981)).

It may be asked why we do not treat the word 'move' in the same way as X, Y and Z, i.e. as a variable and not a constant. There are number of reasons. The most basic one is that the semantic enquiry codes are intimately related to the meaning of the key word such as move: they indicate the semantic role of the variables. Thus, the interpretation rule needs in any case to depend explicitly on the key word. In addition, Dik's researches show that the domains over which the rules apply are also dependent upon the key word. As an example, in a sentence such as 'The meeting lasted for three hours' the variable denoting duration is obligatory (because the verb 'to last' is used), whereas it is not obligatory in a sentence such as 'They discussed the matter for three hours'. Both these facts indicate that the key word should be explicit in the specification of the rules. We assume, then, that the system which knows the rule is one which contains a mechanism which can carry out the process or transformation indicated. The next points to consider are how the mechanism gets its input and what triggers it off. By definition, the process ought to be triggered off if a signal is present that is within its domain and if the conditions relevant to the rule are satisfied. Clearly pattern detectors must first operate to see if patterns falling within the domain of a given rule are present. The information must be held in a suitable buffer, as in the Marcus (1980) model of parsing. Codes representing information under examination must be held by a mechanism which preserves the (in linguistic processing) all-important information concerning original temporal order, and which presents to view the information which is immediately required, concerning the constituents of the buffer.

We know (e.g. from transformational grammar) that the general basis for recognising forms is that of constituent category (plus temporal order). The question of how categories and patterns are actually learnt will be left till later, and we shall consider for the moment the mechanics of transference from the buffer. A significant clue is provided by consideration of related forms such as the active and passive forms of the same utterance. It is normally assumed that the semantic representations of two such utterances are identical (apart perhaps from differentiating markers), representing the fact that their meanings are known to be identical. It is unreasonable to suppose that each of the very large number of semantic interpretation rules present in a given language should be duplicated, while the existence of a transformation mechanism which shunts information around into a different order seems inefficient (though valid as a piece of linguists' apparatus). Instead we can make contact with Dik's theory and postulate an intermediate *labelling process* applied to those elements in the buffer which belong to a pattern which is going to be subsequently interpreted.

This will be illustrated with the example of the two related sentences

$$\text{John kicked the ball} \tag{1}$$

$$\text{The ball was kicked by John} \tag{2}$$

We suppose that preliminary processes recognise 'John', 'kicked', 'the ball' and 'was kicked' (in example (2)) as units and assign labels to them, which are held in the Marcus buffer. The relationship between 'kicked' and 'was kicked' is acknowledged by their sharing a code $\lambda(kick)$, and their being distinguished by features, which include markers for active and passive respectively. The codes in the buffer include information about their grammatical categories (e.g. NP, VP). This information, together with that of the order in the buffer, is sufficient to indicate the different 'presemantic functions' of the units, viz. verb, agent and goal in Dik's notation (we use the word *presemantic* in preference to Dik's *semantic*, since on our model these attributions are not precise determiners of semantic role (instead, the meaning is generated by the interpretation process/rule operating on the presemantic marker)). These presemantic labels (presemantic enquiry codes) are attached to the items in the buffer and used to extract the information. Two modes are possible: immediate interpretation, on the basis of a known interpretation process, or delayed interpretation, in which case a label is attached to the uninterpreted codes, which are stored away using presemantic codes as markers, so that an interpretation process can operate on them later. In the first case, the information in its semantic form can be assessed immediately, so that if necessary an alternative pattern or interpretation can be tried. When a plausible interpretation has been arrived at the information can be stored in a block of memory indexed by a label.

We pause to consider how the Marcus buffer is operated. A suggestive possibility is that the mechanism of *rehearsal* is used. The contents of the buffer normally decay with time, but may be kept topped up by a process which repeatedly scans the contents of the buffer in an order which represents the associated time sequence. Items are no longer maintained in the buffer once they have been assembled into a group, and the code for the group may be added into the buffer in the corresponding place. New items derived, by a unit-identification process from acoustic input, are inserted at the end of the buffer. (A particular implementation of this model, based on labelling the contents of the buffer with ordinal numerals, allows the contents of the buffer to be moved continually towards the beginning as space becomes available).

What happens to the resultant information, which has now been interpreted, or merely grouped on the basis of pattern discovery? This is a matter of the pragmatics of the use of linguistic information. One case is that in which the unit that has been processed is part of a larger pattern (for example either of the above utterances might have been preceded by the words 'I know that'. The nature of natural language is such that the correct procedure is simply to replace the information in the buffer by a code representing the whole (as in the Marcus parser). This is implied by the fact that a unit such as a phrase can be treated as a single entity, with a particular meaning of its own.

The alternative is that the unit is not embedded in a larger construction, and thus functions as a complete whole. It may, for example, be a request such as 'come here', or a declaration such as 'The Moon will be full tonight'. The label giving access to the meaning of the utterance tags a valuable(?) piece of information, and the listener must determine what to do with the label (equivalently, what to do with the information). This question goes under the general heading of pragmatics. As we see from these two examples, pragmatic action is sometimes a matter of social or linguistic convention, and sometimes a matter of general intelligence (which is more or less outside the terms of reference of the current theory). In the first example the convention is to treat the action described as one which is to be carried out, if not overruled by other factors (this raises a general point to the effect that one aspect of pragmatics is to decide the validity, reliability or suitability for use of information supplied by a speaker, and to act upon the outcome of the assessment in a manner which is appropriate). In the second case, the response is not a matter of convention, but one of determining the manner in which the information supplied might be useful, and storing the outcome of any deductions or decisions made, so that similar operations may be tried out in the future.

6 Rules and the goal-seeking aspect of language

What is the actual status of rules in the processing of language? The ultimate purpose of language is to communicate information, and linguistic conventions have a value only to the degree to which they support this aim. We can describe the situation in the following manner: speakers and listeners have various goals, and attempt to achieve them either by applying the processes which correspond to the accepted conventions, or by any other means which may be discovered. The conventions are slaves and not masters. A number of points follow from this. Firstly, it has of some advantages for both participants to conform to the conventions, because the conventions reduce the degree of possible confusion and error. But there are many situations where goal oriented search processes allow the appropriate outcome to be attained without adherence to the conventions. A simple instance is the situation where one component of a group lying within the domain of an interpretation is sufficient to determine which interpretation process is involved (e.g. if the word is a form of the verb 'move', as in the case considered above). The listener may be able to find the remaining components of the group just by scanning the Marcus buffer, conformation to a conventional pattern being unnecessary. The reader will note overlap but not identity with a theme of Wilks (1973), to the effect that syntax plays a subsidiary role in language interpretation, as much of the structure is determinable from the meaning alone.

It is largely the fact that rules can be transcended that gives language its attributes of flexibility and adaptability. The connection between words and meaning is a non-rigid one, the listener's intelligence discovering what is not explicitly indicated by the rules. This flexibility allows the speaker to create new rules, which allow new configurations of ideas to be linked up to language. Furthermore, the fact that a listener has the means on occasion to interpret an item of language even when he does not know the relevant rules is the principal means by which he is able actually to discover rules.

7 New and Old Information, and the Problem of Anaphora

We have so far not taken account of the fact that language may make reference to existing cognitive structures and add to them, instead of just creating new cognitive structures. This fact is taken into account in our theory in a simple way, by invoking particular pragmatic processes. Note firstly that speech may, by virtue of the pragmatics involved, direct the listener's attention to external objects. For example, the utterance 'That's an interesting insect', may activate those parts of the listener's perceptual mechanisms which are involved in attending selectively to insects. In just the same manner, a phrase such as 'the person I introduced you to at lunch yesterday' may activate internal *memory* search mechanisms and activate a code for the particular person denoted by the phrase. In such a case, it is appropriate for this new code to be inserted into the Marcus buffer (or treated by an alternative pragmatic process) instead of the code which corresponds to the relevant phrase. The overall rule is that pragmatically indicated post-semantic processes should be carried out immediately where feasible. We shall not go in detail into the variants which are possible.

Clearly anaphora (the use of part of an utterance to refer to another part, by use of a pronoun or the repetition of a word) is a related problem. The fact that anaphora works efficiently indicates that the method of recording an utterance in memory includes processes designed for easy subsequent recall by anaphoric reference — what may be called anticipatory memory. Codes which it is known from experience may have to be recalled soon are stored temporarily in registers indexed by the words to which the codes refer (so that they are recalled if the word is repeated) or by features such as number and gender so that they may be recalled by subsequent use of a pronoun bearing the same features. Only one item may be stored in a given register, and the pragmatics of the language presumably determine whether a given item is displaced from a register or not.

8 Miscellaneous issues

We shall now discuss briefly a number of points, chiefly concerned with the detection of patterns, which would be relevant to any detailed exploration of the concepts presented here.

(i) It is clearly most efficient for pattern detection to be carried out by a number of independent processors operating in parallel. This remark is subject to certain provisos, however. As already discussed, there are cases where one rule is discovered as a variant of another. In such cases systemic networks of processes are involved, and different patterns may share processing elements.

(ii) The fact that only one pattern out of a number of possibilities is chosen as the basis for further processing indicates the use of the mechanism of reciprocal inhibition.

(iii) We can (to a limited degree) recover from errors by backtracking. This requires mechanisms to remember intermediate states in processing and for effectively erasing errors.

(iv) Patterns may be detected, in principle, either serially (using the original input, or the output of the rehearsal process) or by a parallel process using logic circuitry. Probably both methods are used to some extent.

9 The acquisition of rules

We now come to a crucial issue: how is the complex set of skills that has been described acquired? It is proposed that the fundamental mechanisms are encompassed by the following ideas, which are reproduced, with slight amplification, from the lecture of Maharishi Mahesh Yogi (1974) entitled ‘The Science of Creative Intelligence and Speech’.

We begin with the idea that, for the language learner, order is gradually discovered among what is initially patternless random noise, as the cognitive structures able to represent the significance and form of language develop. At the beginning, speech is represented only at the level of *sound*. As correlations between sound and representations of the world are discovered, response becomes possible to language represented at the level of its sound.

Next, attention is given to the *patterns* visible in language, and as the meaningful patterns become familiar, the processes for detecting patterns and using them are acquired. Thus knowledge of grammar is acquired.

When it becomes possible to pick out the patterns in speech (whose meanings in general are unknown), attention is given to finding out the meanings in other ways, and the diverse rules which a language employs for relating words to meaning are discovered.

Finally, when rules for literal meanings have become familiar, attention goes to discovering what speakers actually intend to imply (often going beyond the literal meaning) by what they say, and thus are learnt the pragmatic aspects of a language (which to a greater extent than the semantics are the shared property of a social group).

(Comment: the above description is not a reference to four separated stages of development. All that is implied is the manner in which acquisition of knowledge at one level is dependent on knowledge gained at previous levels).

Having given this general outline, we proceed to describe the detailed mechanisms. We shall not attempt any general formulation, but discuss instead a number of examples which will illustrate the mechanisms involved.

In the case of familiar objects, such as tables or balls, or categories such as food, or known people, or common actions such as walking or eating, it may be expected that the listener possesses already codes which represent them. Excitation of such codes is significantly correlated with the presence of the corresponding words. We invoke a mechanism which couples receptors responding selectively to the sounds with the corresponding semantic codes. Thus is acquired a primitive type of linguistic rule, one which relates *sound* and meaning. The supposed coupling means that a representation of the object, etc., can be invoked by the hearing of the word, and corresponding responses will follow. To the extent that such outcomes produce positive results (especially of the nature of a concordance between expectation and reality), the rule is reinforced (in the terms of reference of the previous discussion, the conditions of validity of the rule are being established).

The next stage is the discovery of patterns. The only patterns which are useful to know are those with specific interpretations, which convey more information as a group than is conveyed altogether by the elements of the pattern taken separately. We assume that it is on this basis that the patterns in language are discovered.

A pattern is an abstraction artificially imposed on a structure: for any given pattern there exists a defining process which determines whether or not the pattern is present. We therefore propose (bearing in mind the linguistic evidence concerning the nature of linguistic patterns) that the following general procedure is followed:

(i) a meaningful constellation is discovered (according to procedures discussed later);

(ii) a pattern-defining rule which applies to the given constellation is created;

(iii) the rule is applied to subsequent input, etc., to look for more instances of patterns conforming to the same rule.

(iv) these new instances are examined to see to what extent they fit the purpose for which (by virtue of the way the parts of the system work together to perform an overall function) they are intended.

(v) rule modifications are carried out (with the aim of improving the overall utility, in terms of increasing the number of positive outcomes and reducing the number of negative outcomes, of the rule).

(vi) pattern-defining rules generally involve categories (such as nouns or auxiliaries). The categories come into existence by a process of expansion from a prototypical example, as will be explained in due course.

This procedure will first be discussed for the case of standard units of input (such as morphemes, words or phrases). This is a degenerate case of a pattern, in which all examples of the pattern are closely similar to each other (e.g. all the pronunciations of a particular word). We suppose that the system has already learnt to respond to the word, etc. on the basis of its sound. This means that detectors which respond selectively to the sound are linked to codes which function as the meaning. The system does not yet represent the word, etc. as a unit — in fact it has no code to represent the unit as a unit of speech. The change comes about when it does tentatively assign a code for the relevant segment of input. At the time that it does this it must store certain correlations which will function as processes or rules: one which associates features of the acoustic input with the new code, and one which associates the new code with the semantic code (the two associations can be learnt at different times, the second after the first). As hinted at previously, the system has now split its original rule up into two. But the first one incorporates the feature of checking that the input is in agreement with that recorded previously, when the new code was first assigned.

The new code for the word, etc. has the purpose, in a larger context, of indicating that the word has a particular meaning. We now invoke the principle that the system tries varying the rule which defines whether the new code should be assigned to given input, noting whether the input is still associated with the expected meaning. This leads to the system defining the limits over which the pronunciation of the word may range (which may be context dependent, e.g. dependent on accent). We end up with a system capable of responding selectively to the given word. If the word has more than one meaning, it can learn a new semantic rule which links the new meaning with the previous code for the word, without having to relearn how to identify the word.

A similar argument, which will not be expounded in detail, shows that these ideas are applicable to the perception of phonemes in speech. In the above exposition we have implicitly talked in terms of phones (units of sound). The system can in principle learn, by means similar to those already discussed, to interpolate a phoneme detection stage, but here the process which is split in two is the perception of input rather than its interpretation. The hardware for this is almost certainly innate, as is possibly also hardware for working with syllables (the hardware consisting of special memory units, signal channels and so on).

We now come to the subtler problem of learning to detect syntactic patterns. Our theory presupposes that the first stage is that of discovering a *meaningful constellation of units*. The Marcus buffer is a natural tool for doing this.

We suppose that the system has already learnt to recognise a number of common words, knows their meanings, and has also assigned codes to them, which can be held in the Marcus buffer. When a sentence referring to a current scene is heard, codes for some of the words are held in the buffer. The language learner examines his environment looking for something corresponding to one of the codes in the buffer, focusses attention on this part of the environment and then looks again in the buffer to discover further codes corresponding to this part of the scene. In this way it tends to find a meaningful group, e.g. one code may be the name for an object and the other a code for its colour.

The next stage is that changes occur in the system of a type which leads ultimately both to pattern recognition abilities and to the acquisition of an interpretation process for a *group*. A number of developments, which need not all be done at the same time, are needed. The processes will be illustrated by means of a particular concrete example. Suppose the listener hears a sentence (e.g. 'bring the red ball') containing the words 'red' and 'ball' in the order indicated with (for reasons which will become clear) no intervening words in between. Suppose also that a red ball is visible. The buffer contains codes for 'red' and 'ball' (λ_1 and λ_2, say), and by the processes already indicated, the system is able to pick out this pair of codes as being a meaningful group. On the non-linguistic, or semantic side, the system is assumed to have the ability to represent certain facts about the corresponding reality, e.g. that the ball is an object and red is a colour, for which it uses enquiry codes σ_1 (for colour) and σ_2 (for object). Suppose now that the two words of the group are also assigned arbitrary codes π_1, π_2, which will later function as *presemantic* codes. The system also notes the crucial *ordering* of the special elements of the buffer. Tentatively, we hypothesise that codes representing ordinal numbers are used, so that 'red' gets the code 'first' and 'ball' 'second', provided that no interstitial words are present.

The system can now acquire a number of processes/rules by noting the correlations present in the example. At first sight, the codes for red, first, colour and the presemantic code σ_1 should all be correlated together to act as potential process, but a little thought indicates that the actual role of ordinal information in natural language is to provide a constraint for a rule to apply (e.g. for a given type of pattern to be presumed to be present). The net effect of all these processes is to allow the combination 'red ball' to be translated into the correct semantic form (including the enquiry codes), going via the intermediate stage of the presemantic codes. The next stage is for the system to generalise its processes to work with categories of words, instead of with fixed words as in this example. Recall the treatment given for sentences of the type 'move X to Y, using Z' (though here we are dealing with a simpler kind of construction). The new process involves assigning and labelling the particular occurrence of one of the words with an arbitrary label, and treating this label as a variable to be re-recorded in semantic form under the relevant enquiry codes, under a new (generalised) version of the interpretation process (of the form previously described). We want the system to be able to make the correct semantic representation of all groups of the type 'red X'. The hypothesised mechanism for this is that the system collects under a particular category label all instances for which the given interpretation process works, and regards this label as an indication of the domain for which the rule is to be applied.

The procedure just described constitutes the first procedure for category expansion, and is based on essentially semantic considerations: all the elements of the category have similar semantics. If the criterion of similar semantics is dropped, expansion is possible to larger categories, ending up ultimately with the actual linguistic categories of the language concerned. The fact about language that is relevant here is that the rules of a given language which define which types of patterns are meaningful observe syntactic categories rather than semantic ones. For example, the fact that in English the sequence *preposition – noun phrase* is a grammatical combination manifests itself in the two phrases 'under the bridge' and 'under the circumstances' whose semantics are totally unrelated. The implication of this fact is that a procedure of tentatively expanding a category by adding a new word which would fit into an already known pattern if it were placed in that category will often pick up a meaningful combination (though it may also pick up nonsense combinations instead by picking on an inappropriate set to group together). We are led to hypothesise that the system finds tentative category elements in this way, and confirms or disconfirms them by attempting to assign a meaning to the group. This gives a general procedure by which grammatical categories can be established (any erroneous assignments picked up by the means described will tend ultimately to be extinguished, by virtue of the fact that such groupings cannot in general be interpreted by any reasonable rule).

The above process of category discovery, combined with that of pattern discovery, is a cumulative one. The process starts with a few very basic combinations, such as the adjective–noun one discussed above, which lead to the discovery of a few basic categories. These original categories allow new patterns to be discovered, leading to the acquisition of more and more complex patterns, and also those additional categories which apply to particular patterns. It is worth noting, though we shall not go into the matter in detail here, that the gradual growth of a given language probably proceeds through similar pathways.

One of the known regularities of language is that particular patterns generally belong to particular categories (the pattern being regarded as a single entity). For example, in English the combination $ADJ\ NP$ belongs to the category NP. Such relevant facts about the language can be discovered by means similar to those already described.

We note briefly the application of our ideas to a problem mentioned earlier, that of having the system know that, for example, active and passive versions of a given sentence have a similar meaning. Let us take as an example the sentences given previously:

John kicked the ball

The ball was kicked by John

In the first sentence, the pattern *noun – verb – noun phrase* is recognised and the applicable rule for presemantic labelling generates the labels agent, verb and goal respectively (the relevant rule giving the correspondences 1st $\rightarrow$ agent, etc.). The interpretation rules link these presemantic codes to the semantic codes which define the actual semantic roles. Now we hypothesise that a person hearing the second construction but not yet familiar with it attempts to assign presemantic markers to the items in the Marcus buffer in such a way that the previous interpretation rule will still apply. This leads to the last item of the group now being characterised as agent instead of the first. The new rule relating order and presemantic codes is learnt. Investigation of its domain of validity leads to the correlation with the use of the passive form being discovered, and to the specification of this new category.

10 Acquisition of complex semantics

On the basis of its knowledge of the patterns of a language the system can learn the very large number of rules or conventions for interpretation. That these are essentially conventional can be seen from examples such as the one given, 'under the circumstances'. What clearly has happened is that a new phrase has been invented as a shorthand, and has been chosen to conform to a preexisting pattern, partly because of ease of production, and partly because of ease (for the listener) of pattern detection. What the listener has to do to learn to use such unfamiliar constructions is to store the unfamiliar linguistic piece as a labelled group (within the representation of the larger expression in which it is embedded), and also search, by any means at his disposal, for a representation of the meaning. In principle, a cross-correlation process will pick up the interpretation rule, but we have not attempted to examine the question in detail.

11 Acquisition of pragmatics

The principle involved is the same, except that the end product is not the assignation of meaning to a piece of language but the determination of the knowledge that ought to be inferred from the meaning. The dividing line between semantics and pragmatics is probably that semantics involves special linguistic systems devoted to recoding linguistic information into the form of inner speech, while pragmatics involves general inference mechanisms.

12 Innate knowledge

We now change the subject, to discuss the view to which our theory leads us concerning the role of innate knowledge. This view is intermediate to those of Chomsky and Piaget. In our theory we have invoked a number of very specific mechanisms (basically special memories, information channels and correlation detectors). It is reasonable to suppose that these are mediated by innate structures (which if one is so disposed one can speak of in terms of innate knowledge of language). While one might in principle imagine ways in which the 'knowledge' could be *learnt*, such mechanisms would be unlikely to function with the reliability that they actually do have (puzzles of apparently similar difficulty people have trouble with).

On the other hand, we do not need to invoke in our theory Chomsky's 'core grammar'. We have already indicated how knowledge of simple grammatical categories and forms could be acquired.

13 Coordination of component subprocesses

We have described in this paper very many subprocesses which involved in language. The problem of how these are all coordinated together is discussed in the lectures of Maharishi Mahesh Yogi (1974), and in a report by Josephson (1982).

14 Miscellaneous problems dealt with by the theory

Here we wish to illustrate the theory in action by indicating how a number of specific problems are dealt with.

(i) Words with no concrete reference.

In some theories of language it is supposed that words have meanings by virtue of having specific referents, a proposition which encounters considerable difficulties if one tries to implement it. In the present theory, meaning is in the first instance *subjective*, and the link which is encompassed by the word meaning is one which connects together *internal* signals in the system. Thus the domain of meaning is that of the internal language of the system, and any codes it uses can be associated with speech units. In this way we can understand readily how words such as 'not', 'easy', 'round', etc. acquire meaning in their users' linguistic systems.

(ii) The difference between children's and adults' understanding of words (question raised by K. Sparck-Jones).

A typical example is the word 'dictionary'. A child on first contact with the word can have no concept of its actual meaning, but nevertheless can recognise easily a particular dictionary and assign a code to it. In our theory the change which occurs when a child learns the real meaning of the word is an instance of updating of a rule. The original rule which assigned to the word the code for the particular book is changed to a rule which assigns a better code to the word, one which represents its real meaning.

This is not the only type of change possible. In a situation where the word is understood differently because more is learnt about the entity to which the word is applied, the appropriate change is that of storing more information away in the block of memory to which the code for the meaning gives access.

(iii) Language used as a secret code (objection raised by grant-awarding committee — the illustrative example is ours).

A foreign agent spots the spy he is due to meet, approaches him and remarks 'We last met in Pittsburgh, I believe'. The spy replies 'No, it was in Berlin in 1944' (the exchange being a prearranged one to establish identities). The explanation of this interesting piece of linguistic behaviour clearly requires taking into account particular preceding instructions, which we may represent symbolically as follows: the spy is told 'your contact will say A, and then you will reply B'. Without going into full details we can see that the correct pragmatic effect of interpreting this information is to set up a new interpretation rule, which involves treating the sentence A as a special one, having its own individual code a, and assigning a meaning code a' to it, which gives access to a block of memory in which is stored information such as the fact that the speaker is a foreign agent and that the spy should answer with B. (Psychological arousal mechanisms can be invoked to account for the fact that the normal response (treating the statement at its face value) is inhibited).

References

[1] Dik, S.C. (1978), *Functional Grammar*, North-Holland, Oxford.

[2] Eckblad, G. (1981), *Scheme Theory*, Academic Press, London.

[3] Josephson, B.D. (1982), *On Target*, Cavendish Laboratory Progress Report TCM/29/1982.

[4] Mahesh Yogi, Maharishi (1974), numerous videotaped lectures, in particular the lecture course entitled *The Science of Creative Intelligence*, lecture 25 of which is concerned specifically with speech.

[5] Marcus, M. P. (1980) *A theory of syntactic recognition for natural language*, M.I.T. Press, Cambridge, Mass. and London.

[6] Skemp, R.R., (1979), *Intelligence, Learning and Action*, Wiley, Chichester and New York.

[7] Wilks, Y. (1973), *An Artificial Intelligence Approach to Machine Translation*, in Computer Models of Thought and Language, ed. R.C. Schank and K.M. Colby, Freeman, San Francisco.