DOCUMENT RESUME

ED 122 591 FL 006 673

AUTHOR Taylor, Martin
TITLE Speculations on Bilingualism and the Cognitive Network. Working Papers in Bilingualism, No. 2.
INSTITUTION Ontario Inst. for Studies in Education, Toronto. Bilingual Education Project.
PUB DATE Mar 74
NOTE 54p.; For related documents, see FL 006 670-674

EDRS PRICE MF-$0.83 HC-$3.50 Plus Postage.
DESCRIPTORS *Bilingualism; Child Language; *Cognitive Processes; *Concept Formation; Language; Language Development; Language Patterns; Language Research; Perceptual Development; *Psycholinguistics; Second Language Learning; Syntax
IDENTIFIERS *Cognitive Network

ABSTRACT

In the first part of the paper the idea of the cognitive network is developed. The network consists of concepts linked together by relationships which are themselves concepts. Concepts are learned according to simple rules, and the network grows as new concepts are learned. Part II considers the growth and structure of language. The growth of language within the network follows the same rules as the growth of perceptual ability, Labels are attached to some concepts for syntactic transformations to others. Perceptual syntactic relations tend to have syntactic programs as their expression in language. In Part III some problems of bilingualism are considered. The simultaneous growth of two languages presents special problems to an infant. Instead of linking labels and syntactic programs directly to concepts in the network, the infant's linkages must be conditional on extraneous factors. The bilingual infant should have early difficulty with language, but should eventually derive a richer concept structure than a monolingual. Second language learning is seen as a process of breaking down structures that have stabilized in order to replace them with new and more complex structures. Linguistic relativity is real, but occurs primarily in the more abstract realms of thought, not at the perceptual levels. (Author/CLK)

Speculations on Bilingualism and the Cognitive Network

M.M. Taylor

Abstract

In the first part of the paper the idea of the cognitive network is developed. The network consists of concepts linked together by relationships which are themselves concepts. Concepts are learned according to simple rules, and the network grows as new concepts are learned. Lower level concepts are stabilized and become parts of patterns which form higher level concepts. Part II considers the growth and structure of language. The growth of language within the network follows the same rules as the growth of perceptual ability. Labels are attached to some concepts, programmes for syntactic transformations to others. Perceptual syntactic relations, such as "up" or "agent", tend to have syntactic programmes as their expression in language. The general function of syntax is to highlight concepts and to ensure that concepts in a discourse are attached to the correct links in the cognitive network structure. In Part III some problems of bilingualism are considered. The growth of two languages at once presents special problems to an infant. Instead of linking labels and syntactic programmes directly to concepts in the network, his linkages must be conditional on extraneous factors, such as the person with whom he is conversing. Accordingly, the bilingual infant should have early difficulty with language, but should eventually derive a richer concept structure than a monolingual. Second language learning is seen as a process of breaking down structure that has stabilized in order to replace it with new and more complex structures. It is suggested that linguistic relativity is real, but occurs primarily in the more abstract realms of thought, not at the perceptual levels.

Speculations on Bilingualism and the Cognitive Network[1]

M.M. Taylor[2]

INTRODUCTION

The business of perception is to discover consistent relationships among the properties of the real world, that of cognition to order and manipulate these relationships, and that of language to distribute them from one person to another. The unity of purpose among perception, cognition and language suggests that they should be closely linked; this paper suggests that the unity extends to a similarity of basic mechanisms. The purpose of the paper is to present a speculative framework, based largely on work in perception and memory, for the development of language in relation to other intellectual capacities, and to explore some implications for language analysis and for bilingualism.

The fundamental idea developed in these pages is that of a stable pattern of properties (e.g., roundness, bounciness, size convenient to the hand, softness) which come to form a "concept" (e.g., ball). Properties may be relationships "up", "inside", "owner", "agent" etc., as well as static properties. The major thesis is that any sufficiently stable and frequently encountered pattern of properties will come to form a concept, and that concepts may serve as building blocks in forming higher-order concepts. The structure which results is called the "cognitive network".

Language develops from, and as part of, the cognitive network. It is not a separate entity requiring rules of its own, but follows the rules which govern perception and motor skills in general. Language is unique, however, in that it provides a method whereby structures within the cognitive network may be displayed publicly. It will be argued that the syntax of natural language depends very largely on the structural relations within the cognitive network, despite the enormous surface variations in syntax among the languages of the world. It will be further argued that, because syntax does depend on the fundamental structure of the cognitive network, therefore the variations in syntax do not cause or signal variations in the ways speakers of different languages perceive the world. On the other hand, variations in the concepts embedded in the vocabularies of different peoples may both signal and cause such variations in perception of the world.

This paper has three main sections: (I) an introduction to the cognitive network, its structure and the way in which it may grow; (II) an extension of the basic cognitive network to specifically linguistic structures, and a consideration of some aspects of syntax; (III) comments on the effects of a second language on the network structures, and implications for second-language learning.

PART I: INTRODUCTION TO THE COGNITIVE NETWORK

The cognitive network is one model of the structure of knowledge. It supposes that there exist individual "concepts", each of which is linked by a set of relationships to other concepts. For example, a "house" might be a concept. It "has" windows, "is" red, "can be lived in" by a person, "is located" beside a street. These relationships may be indicated by placing the concept "house" at a node of the network, and linking it with other nodes by links which represent the various relationships, as in Fig. 1.

```
                 windows       red
                    ↖          ↗
                     has     is
person <--can be lived in by-- house --is located beside--> street
```

Fig. 1.

The cognitive network, as developed here, is an extension of the network formulated by Rumelhart, Lindsay and Norman (1972) and by Rumelhart and Norman (1973). A clear introduction is presented by Lindsay and Norman (1972) in their introductory textbook, and a reading of their chapters 10 and 11 is recommended as an introduction to the present paper. In their formulation, the network contains not only static knowledge patterns, such as the relationships of the concept "house" in Fig. 1, but also programmes for executing action. It contains all the long-term knowledge of the individual, such as how to ride a bicycle, the way the buildings are arranged on his street, Pythagoras' Theorem, the fact that Aunt Martha came to dinner last Tuesday and didn't like the dessert, that one does not normally say "he ain't did gone", how to catch a thrown ball, and all the knowledge that permits us to put some of these ideas into words.

Lindsay and Norman (1972) do not explicitly consider the role of the cognitive network in language processing. However, the ideas of the network were originally influenced by Fillmore's case grammar (Fillmore, 1968), and Rumelhart and Norman (1973) have shown how the knowledge in the network is required to disambiguate ambiguous sentences. Norman (personal communication) indicates that the network concepts are being used in a wide variety of applications, including linguistic theory. The application to language considered here is independent of those extensions, although the ideas may well turn out to be similar.[3]

The gross structure and environment of the cognitive network.

A person's cognitive network must communicate with the outer world. In a sense, there are two outer worlds, the world of things and the world of people. The difference is that people talk, and as we shall see, talking is a technique explicitly designed to pass structure to and from the cognitive network. There are four channels whereby the cognitive network communicates with the outer world: Perception, Motor Actions, Talking and Listening (the latter pair include writing and reading, respectively). The motor and perceptual processes are probably very tightly linked together. Some would go as far as to say that there is no perception without the possibility of some contingent motor action (e.g., J.G. Taylor, 1962). According to this view, the perception is the activation of a engram linking the stimuli with potential response. The talking and listening channels are likewise tightly linked. Some (e.g., Liberman, 1957) have proposed that understanding speech depends on the ability to determine how the speech was produced, or even on the covert production of the same speech. While this strict view is no longer current, "Analysis by Synthesis" ideas (e.g., Neisser, 1967) similarly imply a tight linkage between listening and talking. I do not find the arguments convincing, although some degree of bonding between input and output processes both at the perceptual-motor level and at the linguistic level is almost certainly required for the individual to function effectively. In the following, we shall seldom be directly concerned with the influence of output processes on input processes, and will ordinarily consider all four channels of communication between the network and the outer world as if they were separate.

The network itself, like any network, consists of nodes connected by links. The nodes represent concepts, the links relationships among concepts. Some links represent class membership (a canary "isa" bird), some represent properties (a house "has" windows), some represent modifications of one concept by another (this roof "is" blue). In these examples, "isa", "has", and "is" represent some of the many different types of link.

Activity in the cognitive network consists of a "passage of thought" from concept to concept through the connecting links. If a node is activated, so are other connected nodes, though possibly to a lesser extent. We will not attempt to define this passage of thought any further. It is a conceptual convenience, not a theoretical construct. Lindsay and Norman (1972) use in a similar way the concept of a flashlight beam which illuminates part of the network. But some kind of device for highlighting parts of the network is required to account for many of the phenomena of thought and language. A flow of activity is consistent with a neurological view of the network and fits in with association effects and the flow of language.

Since the network contains programme as well as data, the flow of thought through it produces action and speech as well as mental operations. The programmes are not readily distinguished from the data on which they work. Vocabulary and syntax are both held in the network. Speaking and interpretation are both network functions. If the person is multi-lingual, the network must be more complex than it is for monolingual persons. The differences between multi-lingual and monolingual networks should depend on the manner and the context of learning second and subsequent languages, and we should be able to predict some of these differences from consideration of the way the network may grow.

Structures and growth of the cognitive network.

It is very difficult to talk accurately about the structure of a network. The nodes in the network are not laid out nicely in a plane, like those of a fishnet. Any node may connect to any other. There may be a multiplicity of nodes referring to a single concept, or nodes may have different degrees of salience. Some pairs of modes may be connected by links that are stronger or weaker than other links, or alternatively the apparent variation in associative strength between concepts may depend on the number of links which connect them. None of these subtleties are reflected in the network drawings used in this paper.

The growth of a conceptual structure is very much a question of pattern recognition. If we have a concept, such as a "house" (rectangular solid object, which has walls, a roof, windows, and can be lived in) and we see a rectangular solid object with walls, windows, a roof, and which looks as if someone lives in it, then we consider what we see to be a house. In describing it to someone, we do not mention the specific properties that lead us to call it a house, but rather we assume that they know those properties as soon as we say "house". We describe only the properties which distinguish this particular house from all other houses: it has a green door, a red roof, lace curtains in a bay window, and so forth. These relations can be depicted in network form as in Fig. 2. (Some of the properties are omitted to make the figure clearer).

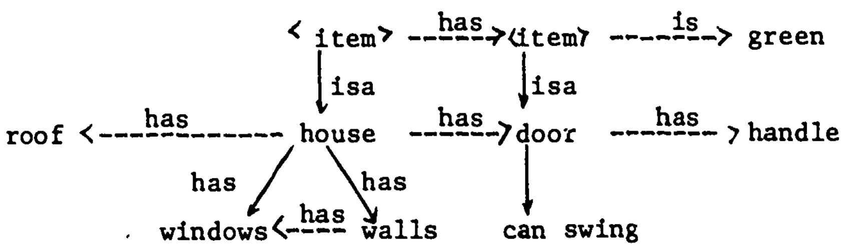

Fig. 2

The <item> nodes represent particular occurrences of something in perception, not general concepts. The item which "isa" house, "has" another <item> which "isa" door. This particular door, but not "door" in general is green. But by virtue of being a door, it "has" a handle, and can swing. The first <item>, by virtue of being a house, has a roof, walls and windows.

The identification process for "house" demanded recognition of the fact that the <item> had a roof, walls and windows, and matched this list of properties with that of the class "house". As soon as this match was made, there was no longer any need to maintain a record of the linkage pattern from the <item> to the "house" properties. The recognition process can be summarized as in Fig. 3. In stage 1, the <item> has not been recognized. By stage 2, it has, and the property linkages which make it a "house" have been deleted from its representation in favour of an "isa" link to <house>.

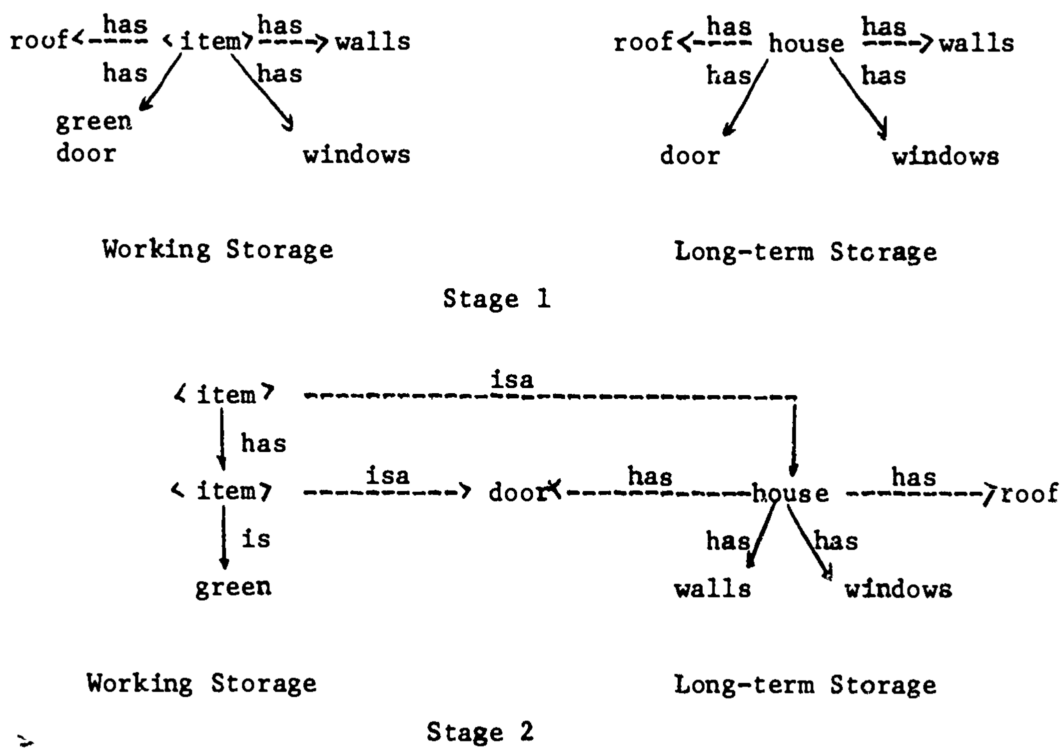

Fig. 3

The recognition process, as sketched, requires that the recogniser has previously recognised walls, roof and door. But the recognition problem does not imply an infinite regress, for two reasons. Firstly, there is the matter of context, of highlighting. Only in laboratory experiments are we confronted with an image of which we must make some sense with no context. Ordinarily, when we see a house, we are in the context of houses; we see roads, paths, other houses, and so forth. Hence, the concept of "house" is highlighted, and the recogniser has only to match the "house" pattern of properties with the property pattern of the <item>. The roof, walls, and door in their turn are recognised as much because they fit into the house pattern as because of their own distinctive features.

The second reason that the process is not an infinite regress is that each concept is eventually defined in terms of directly sensory processes, edges, corners, lines, colour patches, and so forth. These are unique and probably very little dependent on context. According to most theories of pattern recognition, these most peripheral patterns are coded in terms of features which suffice to describe the input patterns. Features are the simplest concepts in the network. A pattern which more or less matches some feature activates that feature to some degree. Most patterns will activate many features, and the degrees of activation of the various features define the pattern. Under certain assumptions (Taylor, 1973) features will evolve so that the network of peripheral features will come to be an efficient description of the input.

Patterns activate features. These features themselves form patterns and enter into relationships, which can excite higher-level concepts; this process may continue until the activity reaches the level of concepts like "roof", "house" and so forth. The process depicted in Fig. 3 is one such stage. One should note carefully, though, that identification of the pattern as a "house" in no way prevents its simultaneous identification as something else, such as a "pile of bricks" or a "reinforced concrete structure". Most inputs excite many features, and are perceived as having many properties. One item may have "isa" links with many concepts. Different properties of the item are effective in arousing different concepts. The chair in the living room may be a "chair" but it may also be a "red thing", or a "soft thing", or a "grandfather's thing", or a "very valuable thing", depending on what aspect of the world is salient at the moment.

Concepts do not represent only object properties. A concept comes into being because a particular pattern of relationships recurs. Indeed, the notion of a particular relationship is itself a concept, which has grown because many pairs of objects were related in that particular way. Relationships such as "above", "inside", "behind", are quite elementary relationships, which require only the segregation of objects in the world for their definition. A relationship such as "part" depends on objects being viewed as constructed from other objects, and may be more sophisticated. A relationship such as "lives in" depends on a great deal of prior conceptual structure. But all relationships can be dealt with as concepts built because they recur in experience, and can be used as components in high-level concepts.

The human is very good at pattern matching, as compared to any automatic algorithm so far developed. It is probably not too much to say that pattern matching is largely what the brain is specialized for. We see consistencies and relationships where none even exists, but may discard them when the pattern does not recur. The repetition of patterns is detected, often even though no concept for that pattern exists in long-term storage. If the repetition recurs, a concept is added to the network, and another recurrence "isa" example of the new concept. But notice carefully that the elements of the new concept must be concepts which already existed in the network. Patterns of properties seldom, perhaps never, grow in isolation, although teachers sometimes try to make them do so.

Stability and Critical periods

The existence of critical periods for learning different aspects of the world seems likely. Children readily acquire a good accent in a new language, whereas most adults have great difficulty in doing so. "Wolf" and "attic" children have trouble learning language at all, if they have not been exposed to it when they were young enough. At a more peripheral level, Blakemore and Cooper (1970) found that they could permanently fix the distribution of the orientations of line-detector cells in kittens by exposing them for a matter of a few hours at the correct age to an environment having only horizontal or only vertical lines. Pettigrew and Freeman (1973) exposed kittens to an environment full of star-like spots but no lines. Their kittens developed no line detectors, just spot detectors. There seems to be a period during which the line detectors develop. If they do not develop properly during this short period, they never do.

A succession of critical periods is almost a requirement for a developing hierarchical system like the cognitive network. A higher level cannot be securely built unless the concepts on which it is based have become stabilized. The structure is like a building, which cannot stand except on secure foundations. Humans, however, seem to be more adaptable in most respects than lower animals, and the stabilization may not be as rigid as suggested here.

The mechanism whereby a concept level is stabilised is not clear. One suggestion (Taylor, 1973) is that the concepts at a particular level interact by lateral inhibition in such a way as to distribute them as widely as possible over the patterns which actually occur in the input to that level. The implication of this idea is that the prominent, most common patterns in the input are first developed into concepts. The cells which serve these concepts then inhibit other cells from responding to the same or highly related patterns, so that these other cells come to respond to less prominent patterns. These in turn inhibit the cells as yet uncommitted, so that very soon, all possible patterns are served. The prominent patterns are served by many cells, the less frequent by fewer. Furthermore, the structure is fairly rigid, in that a cell cannot change its committment very easily, because it is inhibited by other cells when it might otherwise respond to a slightly different pattern from the one to which it is best tuned. The level is stabilized almost as soon as it is grown. It may well be that a concept level will be destabilized only if the input patterns change to include patterns which cannot be described in terms of the original features of the level. In this case, the lateral inhibition among the concepts may tend to "drive" some cells to respond to the strange pattern.

Another mechanism which stabilizes a concept may be its use as a feature in the pattern of a high level concept. Feedback between a concept and its properties is a property of the network. Lower level concepts may be stabilized just because they are at a low level.

The pattern of phonemic categories is a good example of a stabilized concept level. One of the most difficult problems for an adult student of a new language is to hear phonemic distinctions between sounds he has learned to categorise together, and to ignore distinctions between sounds he has learned to call different. It is even more difficult for him to learn to produce the sounds correctly under such circumstances, a problem at the same level of the network, on the output side. If the question of learning phonemic categories is seen in this light, it seems clear that children should be exposed to a wide variety of phonemic systems during the phonemic critical period, even though they do not learn the relevant languages at that time, so that later they may be able easily to hear the distinctions when they do want to learn the languages.

Data and Concepts.

Concepts in the network grow from recurrent data. The long-term memory retains both the concepts and at least some instances of the original data. If I ask you what you did last Tuesday, you may be able to give me a sequence of things, such as that you got out of bed, had breakfast, caught a bus to work, and so forth; all things which you normally do. But then I say, "But I asked you about last Tuesday?", and you think for a while, and then report that you felt a bit ill in the morning and so stayed in bed until you felt better, then decided to take the rest of the day off and went for a short walk to get some fresh air, and so forth. The individual sequence of events was stored, and dated. The first answer ran off the sequence of events which together had come to be stored as the concept "what I do on a weekday". The second found the specifically dated string of events which happened once only.

Now suppose that for a few weeks you always felt a bit ill on Tuesday morning, so that the pattern of events which had been restricted to the one accasion became a Tuesday routine. Then, if I ask you what you did last Tuesday, the new routine will be reported without hesitation. The idiosyncratic event, without being lost as an individual event, has become a routine. But once a pattern has become a concept, its individual elements may not be noticed except insofar as they confirm the pattern. How often does one drive over a familiar route and immediately afterwards remember nothing of the drive except that it occurred and was routine?

This example points up another rule of cognitive network structure. Generally, the components of a pattern are less likely to be noticed than is the total concept. The highest level of abstraction relevant for current action is the level with the most prominent activity. Rumelhart and Norman (1973) put this notion forward as a problem. Paraphrasing Abelson and Reich (1969) they quote the problem of the three drugstores:

Statement: I went to three drugstores.

Response 1: How did you go?

Response 2: What did you buy?

Response 3: Why did the first two drugstores not have what you wanted?

These three responses are ordered by the conceptual distance between the statement itself and the problem raised by the statement. The first response depends directly on the occurrence of "went" in the statement. The listener knows that "went" can take an instrumental case, which can be filled by a phrase such as "by car", "on foot" and so forth. The second response depends on some knowledge of the world. Part of the pattern invoked by going to a drugstore is buying something; the question is what was bought. The third is even more abstract, depending on the knowledge that one would ordinarily buy at the first drugstore all the drugstore items one wanted; but of the three responses it is the most likely for a human to make, provided that neither of the other levels was currently highlighted. The statement about drugstores arouses many patterns of activity in the cognitive network. Some of these patterns invoke open links, links whose value could in principle be given by the person who made the original statement. The most active of those open links should give rise to the question actually asked by the repondant.

Schank (1972) considers a similar "level of abstraction" problem in dealing with our prediction of what is about to be said by someone in the middle of an uncompleted utterance. He identifies six different types of prediction (paraphrased here):

(1) what kind of syntactic category is likely to occur;
(2) what conceptual category is required to fit into the simplest (syntactic) conceptual diagram;
(3) what information would fit into the context of the conversation;
(4) what would "answer a question" arising from prior context;
(5) what we know about the attitudes of speaker and hearer to the topic;
(6) what will best relate to the total memory structure of the hearer.

These six categories more or less cover the levels of abstraction in the conceptual network. They tend to derive from one another by successive transformations of patterns into more abstract and all-inclusive patterns. The meaning of a sentence occurs at many levels, from literal through psychosocial. Meaning should arise through the arousal of sequences of patterns in the network. Activity flows should diverge and converge on many concepts, all of which convey the meanings of the conversation. Meaning is associative, connotative, and denotative.

In terms of the levels of abstraction, once-only events may be considered as the lowest level. Each component of such an event has a conceptual reality of its own. But groups of these components also combine to form concepts: being ill -----> unpleasant; staying in bed -----> pleasant. Connotative meanings abound. Hence, even an idiosyncratic pattern, one which never recurs, can find its linkages within the network.

Are individual occurrences all stored? This is a matter of much argument. If an individual occurrence is stored, it must surely be stored in terms of its component properties, event sequences, and so forth. It cannot be abstracted as a concept of its own, because a concept is defined in terms of its properties. But if a major part of the same pattern recurs while the original is still sufficiently highlighted, the common components may match and become a concept. Then the original item would be stored, plus the newly generated concept. Maybe the subsequent occurrences would also be stored, at least in terms of their own unique elements. They might be remembered as comprising the concept and other features as well. It is hard to say. In any event, it seems likely that the network contains all sorts of unique information, specialized concepts, and more inclusive and abstract concepts. These latter are the most likely candidates for integration into the language, since they occur often enough to make it worthwhile to provide them with unique words.

Relational concepts.

Not all recurring patterns produce static concepts. Many, perhaps most, are dynamic or involve relationships among static patterns. The concept of someone doing something must be very strongly embedded in the cognitive structure. Most of the time we see people, they are actively engaged in affecting something else. This abstract relationship, which we may call "agent", is extremely common, and should be very early derived as a component of almost all patterns of events involving people or animals. Other relationships are almost as ubiquitous. When someone does something, he usually does it to something. The object of the action is a very common relationship. The action itself, or the concept of action, is also a very frequently occurring relationship among people and other people or things. We can demonstrate a whole abstract pattern of relationships, as follows: someone does something to something for someone with some instrument by some method for some reason. All of these relationships are common, and should occur as concepts in their own right within the cognitive network. When language develops, they should find some means of expression, regardless of which language it is. If universals of language exist, surely they should be found at this level.

Generally, the concepts which define relationships tend to drive action programmes rather than invoke static images. My image of "gave", for example, is a dynamic picture of someone actually passing something over to someone else, not a static picture of "a giving". I have similar dynamic images for the more abstract relationships of "agent" -- a multiplicity of people do all sorts of things; in some sense, "agent" is what is common to all these events.

When it comes to speaking, the relationships tend to drive transformations on the words with which we label concepts. They select word orders, affect number and gender agreements, select propositions and postpositions, change the forms of words by declension and conjugation, and so forth. The more abstract the relationship, the more likely it is to be expressed in a "syntactic" manner, rather than as a content word. These concepts tend to be expressed as programmes rather than as data.

Relationships as nodes in the network.

We have discussed relationships as concepts in their own right, but we have not so far considered how they may be displayed in our network diagrams. Furthermore, the proliferation of relationships of various kinds does not enhance the credibility of the network concept. In this section, the various kinds of linkage will be reduced to two, the "isa" link and a simple directed association.[4] "Isa" indicates that one concept partakes of all the properties of another. A house "isa" dwelling, but not the reverse. If an item "isa" house, it automatically partakes of all the properties of "house" and thereby of the properties of "dwelling".

To see how the other kinds of link may be reduced to simple directed associations, consider the network corresponding to "John gave Jane the book".

$$\begin{tikzcd}
\text{John} & \text{gave} \arrow[l, dashed, "\text{Agent}"'] \arrow[r, dashed, "\text{Recipient}"] \arrow[d, "\text{object}"] & \text{Jane} \\
& \text{book} &
\end{tikzcd}$$

Fig.4.

This simple segment of network has three different types of linkage - "Agent", "object" and "Recipient". If we rewrite the relationships as concepts in their own right, we have Fig. 5.

$$\begin{tikzcd}
\text{John} & \text{Agent} \arrow[l, dashed] & \text{gave} \arrow[l, dashed] \arrow[r, dashed] \arrow[d, "\text{object}"'] & \text{recipient} \arrow[r, dashed] & \text{Jane} \\
& & \text{book} & &
\end{tikzcd}$$

Fig. 5

While this appears on the surface to be not much different from Fig. 4, the names of the links being written between two arrows instead of over the arrows, there really is a fundamental difference in the conceptions of the figures. Firstly, none of the arrows are labelled. They have a head and a tail, and that is the sum total of their properties. But the relational concepts, such as Agent, can now be linked to their own properties, which are programmes:

```
(Output)                              (Input)
<   >-----------Agent  ----------  <  >
                  |
                  v
(i) Conceptual:  The concept at output
                 performs the action at input.
                 (Someone does something.)
(ii) Linguistic: In English, concept at output
                 precedes concept at input if it
                 is the topic; else  concept
                 at output follows concept at
                 input and is preceded by "by".
```

Fig. 6

The properties of such a concept (which are only partially stated in Fig. 6) relate to it in exactly the same way as the properties (has walls, has windows, has roof, has door, can be lived in) related to "house". If the pattern of properties exists in a stimulus, whether verbal or perceptual, then "Agent" is excited. If "Agent" is activated by the passage of thought through the network, then its properties are excited. The unlabelled link in Fig. 6, connecting "Agent" with its properties, is a form of "isa" link, indistinguishable in principle from any other "isa" link.

Network: summary.

At this stage, the cognitive network has been sketched in its mature form and simple growth rules have been specified. Hints have been dropped as to the function of the network in language, but these have not been developed.

The network consists of concepts linked by directed associations and by "isa" links which define a sequence of hierarchical levels among the concepts. Concept A "isa" concept B, if and only if A partakes of all the properties of B. In other words, B is part of the pattern which defines A. Activation of a concept also activates all its "isa" links, highlighting the entire pattern of properties belonging to that concept. The longer the sequence of "isa" links between a concept and a particular property, the less likely is the property to be strongly activated. General properties are less likely to be noticed than are particular properties, but they are to some extent highlighted by activation of the particular concept and are thus liable to be used. Properties may be action programmes as well as data.

The directed association links may not be elementary in neurophysiological terms, but they may be thought of as elementary in terms of the network, since most other relationships may be defined in terms of them and proper conceptual nodes. The actual direction of the arrows in the network is a matter of convention. The growth rules of the network are fairly simple. A concept is generated when a particular pattern of relationships recurs often enough. Concepts may enter into the patterns which define other concepts, thus defining a hierarchy. The activation of a concept also tends to activate related concepts, thus setting up a feedback system whereby concepts which share properties tend to activate one another. However, a system of "lateral inhibition" tends to affect both the growth process and the moment-by-moment activity patterns in the network. Two related concepts that have very similar patterns of properties tend not to be activated together, because the more active tend to inhibit the less active. As a result, some concepts tend to be built for infrequent property patterns rather than all clustering around the most common patterns. Levels of the concept hierarchy tend to become stabilized soon after they begin to be built on the foundations of lower levels. This factor leads to "critical periods" for learning fundamental features such as phonemes, or visual elements such as oriented line detectors.

The four channels of communication between the outer world and the cognitive network are all involved in its development. These channels are (i) perception, (ii) motor behaviour (iii) heard (and read) language (iv) spoken (and written) language. The first two may be regarded as building the network by the growth rules described, while language operates directly on the structure of the network. We shall consider language in the following sections.

PART II: LANGUAGE AND THE COGNITIVE NETWORK

The essence of language is understanding. Structures from one person's cognitive network must be transferred to another's. Hence, we must consider how the structure of language permits this transfer. As Schank (1972) has pointed out, the question is not h w the syntax of the language is constructed, but how concepts are properly related to one another. Schank's model demonstrates that even syntactically very simple sentences may have a quite complex conceptual structure, and for their proper understanding require deep knowledge of the way the world works. On the other hand, the syntactic structure of the language surely helps understanding. We will consider here how the content concepts evolve and how the syntax can operate both in speaking and in understanding language. We will not consider explicitly the deeper questions of understanding such as are exemplified by the following dialogue:

"Hi, Mary. Are you doing anything this evening?"

"Not yet; any ideas?"

(John thinks: What a wonderful world. She likes me.)

We will be more concerned with the questions of surface content, while recognising that the concepts activated by the surface content will necessarily set up waves of action within the cognitive network, interacting with all sorts of structures which may be communal, or may be idiosyncratic to the listener. Some consideration of these effects is often required, even to interpret the surface content correctly, since the speaker may assume that the listener has certain knowledge, and leave that unsaid. For example, "John went to the park with Mary" does not make clear whether they went together to the park, whether Mary awaited John in some particular one of several parks known to speaker and hearer, or whether Mary is a statue in a particular park. This may well be impossible to disambiguate through any amount of prior or subsequent context of the conversation, while being perfectly clear to both speaker and hearer, who have the necessary background knowledge. We shall not consider the questions involved in disambiguating such sentences. Many examples may be found in articles on the interaction of semantics with syntax; Schank (1972), in particular, presents some fine examples, including the above. We shall be concerned with the interpretation of more straightforward material, and with how the natural growth of the cognitive network leads to language production and understanding.

The problems of deep content and social meaning seem more properly to belong to the elaboration of activity within the non-linguistic part of the network than to a theory of how the network deals with language.

The initiation of language.

Most words are labels. They label concepts in the network more or less directly. When a word is heard, the various concepts labelled by that word tend to be activated. If a concept is to be mentioned in speech, the various labels relating to it tend to be activated. Since, as a rule, several levels in the concept hierarchy will be activated by any thought, there will be many ways of saying any particular thing. But any one concept probably is most closely linked to one or two labels. We will simplify the discussion by ignoring the multiplicity of concept levels, and by assuming that in a monolingual, only one label corresponds to each concept.

How does a concept originally acquire its label? Properties which occur together often enough make a pattern which can form a concept. The same is true of simple concrete concepts and their labels. The child has acquired some concept, say <ball> , before he can speak or understand speech. (In this section, a concept denoted by angle brackets, <ball>, indicates the perceptual concept, unlabelled. The label is denoted by quotes, "ball"). Before labelling, the baby's network for ball might look like Fig. 7.

<round> ←---is------ <ball> ------is-----→ <squashy>

<will bounce if thrown> <can hold in hand>

Fig. 7

Often, when the baby plays with his ball, it will be highlighted in his thoughts. His mother may sometimes say "ball" at such times. Eventually, the hearing of "ball" may become one of the properties of "ball", albeit an occasional, or optional one. The network may develop in two phases, so that the label "ball" comes to activate the concept <ball>. In the first phase, there may be two separate concepts, one <ball$_1$>, as shown in Fig. 7, and the other <ball$_2$> which is the same except for the added property <has noise "ball">. According to the growth rules for the network, these two concepts would very soon give rise to a concept whose properties were those which pertained to both original concepts, leaving only the residue of idiosyncratic properties (in this case, the label "ball") to the two individual concepts, as in Fig. 8.

```
                     other properties
                            ↑
round<------is------<ball> ------is------>squashy
              isa  ↗        ↖  isa
        <ball₁>                <ball₂> ----has---->noise "ball"
```

Fig. 8

There is a pattern here. As the baby learns other noises which belong to other conceptual patterns, the general network configuration of Fig. 9 recurs again and again. By the growth rules, this configuration

```
      < xxx > ------------> properties
   isa ↗   ↖ isa
< xxx >     <xxx>  ---has--->noise "     ".
```

Fig. 9

will come to be a concept in its own right, which we may call the "label" relationship between the noise and the pattern of properties which form the base concept. This total relationship can then be rewritten as in Fig. 10, or in the easier shorthand form of Fig. 11.

```
< xxx > ------->properties          <xxx> --------->properties
   |                                  |
label  ------->label programme        | label
   ↓                                  ↓
 "xxx"                              "xxx"
```

Fig. 10 Fig. 11

For the ball, the child then has a piece of network like Fig. 12.

```
                 other properties
                        ↑
round<---------- < ball > ------------>squashy
                      | label
                      ↓
                   "ball"
```

Fig. 12

This may be an extreme example, but it suggests the importance of peer feedback in the development of language. If, for "pik" we substitute "baw", which has been learned by each twin from adult contact, it seems likely that this incorrectly pronounced word will tend to persist. Each twin's use of the word matches that of the other. There will be no perceptual discrepancy between his own output label and the input label he hears most often. He may come to learn "ball" as a separate label, in a different form of speech used in a different social context. Twin speech then comes to be a special form of bilingualism, where the total monolingual community for one of the languages consists of two people.

Since input and output labels tend to coincide after a longer or shorter period, they may conveniently be linked as a single node in the network. But this is only a notational convenience, and tends to obscure the very real difference between a heard word and a spoken one. There is no logical reason why the word spoken by a child should match the word understood by him as a label, and it is a task of theory to describe how the observed coincidence does come about. The approach taken here is to suggest that the child can hear himself, and is able to notice that his productions do not sound like those of other people. He may well not know how to change his productions to make the match. Indeed, he may not even be able to control his speech accurately enough to make the match consistently. In making random variations, however, sometimes he will come closer to the sounds other people make, and perhaps what he does on these occasions will tend to be held as modifications to the output programme.

Children become marvellous mimics soon after they learn to speak reasonably well. The growth rules of the network suggest that this may be so because of the continuous variation of production involved in learning these early output labels and making them match other people's speech. In making these shifts, the child learns concepts involving high-level variation in speech production. The production programmes may contain parameters which can be systematically varied for "more nasality", "darker sound", "flatter intonation contour" and so forth. These would be abstracted from the specific programmes which produce adequate labels in the same way that higher level perceptual concepts are abstracted from recurrences of patterns at lower perceptual levels. Some adults maintain this ability to mimic, while others allow it to lie dormant from the age at which they first achieve what is to them a satisfactory agreement with the speech of their community.

We have discussed the acquisition of single words as labels for input and output. The theory suggests, firstly, that children should learn to begin with that some sounds denote certain concepts, and later learn the relational concept of <label> which permits rapid acquisition of new labels; secondly, that the child should independently learn to make some noises which are accepted by other people as labels in such a way that the relevant concepts become highlighted, followed by a stage in which the concept <output label> is derived so that the child can attempt deliberately to label things; thirdly, that output labels should begin to be identified with input labels, so that the child can discover words by matching output to input and by mimicry. Only after all three stages (of which the first two may be simultaneous) have been completed can the business of language learning begin in earnest.

From word to phrase.

When the child has a command of some labels, he is in a position to express some of his concepts, and to respond to the speech of others. But in either case, one may presume that only those concepts signalled by separable words (or word groups like "come here") will relate to speech. A concept may be activated by a sound pattern only if that pattern has been linked to the concept. Relational concepts such as "agent", even though they may be well formed in the child's cognitive network, should not be expected to be linked to syntactic structure programmes until quite a few labels have been stabilized. Relational concepts arise because the same relation recurs over and over. Syntactic patterns may be directly associated with these relational concepts in the mature network, but in the growing network of the child, they cannot be so associated in the early stages. Only when object labels and action labels have been learned in sufficient numbers can syntactic structures be learned independently of individual phrases. "Come here" is probably treated as a single word before it is related to "come" and to "here". The same thing applies to words of several morphemes. The whole word should be a concept before its morphemic structure is abstracted. Later, of course, morphemes can be used as components of words, and participate in syntactic operations.

The flow of thought through the network activates concepts. The child may utter labels for those concepts he can. The result for a transaction may be one word or several independently uttered words. Bloom (1973) describes a period of some months during which sequential separate words are used in the form of connect, but non-syntactic discourse. The following example is quoted from Bloom (1973, p. 50). Allison is aged 19 months, 2 weeks.

(Allison picks up a blanket; hands it to Mother)	"blanket" "cover"
(Mother says: Blanket? Cover?)	
(Allison touches doll's head)	"head"
(Allison touches doll's head and lifts doll to her own head)	"head"
(Allison touches doll's head in front of her)	"head""head"
(Mother says: Head?)	"cover"

Allison obviously wanted her mother to cover the doll's head with the blanket, and had this whole event in mind, but was able to only label the objects and main action. Quoting Bloom again (1973, p. 52), "Allison's single-word utterances were apparently mapped onto successive movements first, and then, subsequently, successive single-word utterances were mapped onto the mental representation of a whole event (object or relation) in experience." This period lasted from age 16 months to age 21 months, during which time she learned many labels and used them with increasing assurance.

In a later stage of the single-word period, the children studied by Bloom used several relational words (there, no, away, gone, and more are specifically quoted) and these are the very words which later were used in the first sentences. But "it was obvious that they had yet to discover the way in which the language they had been exposed to coded such basic notions in terms of semantic-syntactic relationship among words. One thing was abundantly clear, however. An hypothesis of insufficient vocabulary could not account for the nonoccurrence of sentences. The children certainly used enough words for sentences; in fact, they used precisely the words in succession which they could be expected to combine in two-word utterances eventually." (Bloom, 1973, p. 55).

From the point of view of the cognitive network, this pattern of growth is precisely what should be expected. The children have the concepts, and learn words to suit. But the abstraction of word patterning must await the stabilization of the necessary words. The concept "agent" requires, in one pattern, that the label for the actor precede the label for the action. One must know the label for the actor and for the action, and have the whole relationship highlighted when the adult sentence is heard, before the pattern can have an opportunity to be developed even as an input concept. The pattern "possible actor label followed by action label" must attain status as a concept for input sentences, and must be found to correspond with the "agent" relationship in the concept structure, before much useful learning of output syntax can occur. The situation is precisely parallel with the first learning of labels, discussed above. One may guess that at some point two words which are ordinarily spoken in sequence will be spoken as a pair, and that pair will be heard to match some already learned syntactic pattern.

Syntax.

We have come to the position that early syntax is a matter of pairwise relationships, and that it is learned in a manner exactly parallel to the first vocabulary. The syntactic rules which first apply are those which identify the common relational links within the cognitive network. Furthermore, although syntactic needs may be the same for different languages, the syntactic mechanisms may be entirely different. There is no difference in principle between learning (in English) that an object word which precedes an action word is usually an actor and learning (in Japanese) that an object word followed by "-wa" is probably an actor.

Syntactic relationships are relationships among words. From the viewpoint that the function of language is to transfer structure from one person's cognitive network to another's, the only reason for these relationships to remain consistent is that they should signal consistent patterns within the cognitive network. In other words, all rules of syntax should find their counterpart in aspects of the network structure. If we turn this around, we should find that particular relationships within the cognitive network drive, as labels or programmes activated by those relationships, syntactic function words and syntactic transformations. In listening, word patterns should drive transformation routines which activate the appropriate relational constructions in the network.

This position suggests that the complexity of individual syntactic rules in mature speech should match that of the structures they signal. Most relationships in the network are bilateral. The "agent" links the actor with the act, and so forth. Some relationships are three-way. A verb, such as "gave" may link three items, the agent, the recipient and the object. Stevens (1973) attempted an exhaustive listing of English verb types, finding 25 different types, and discovered none with more than three relationships. Not counting two types, which each had only a single exemplar, he found 16 types with two required relationships, and only 7 with three. A three-way relationship seems to be the most complex allowed in modern English. Schank (1972) also claims three-way relationships to be the most complex that occur in the related concept structures.

It is often said that an English sentence requires a verb. An English utterance does not. A sentence is supposed to be able to stand in isolation. If it is to pass a structure from one person's network to another, it must match that structure in some way. But the only type of relationship which links two object-concepts is one which is signalled by a verb in English. Hence, any structural relationship must include a verb (in English), and an utterance which is to stand alone must incorporate that verb. This is not true of utterances in context, particularly in the feedback context typical of conversation. Writing does not permit feedback and the structure of written langauge is more formal in consequence. In a conversation, each utterance highlights a portion of the network, and leaves links open or potentially open in the structure of the receiver's network. He can request that these links be closed, or attempt to close them, without restating any relational constructs:

"How's your wife, Joe?"

"Much better, thanks."

or again:

"Mary went to the concert with John yesterday."

"Mary? with John?"

"Yes, those two."

The initial utterance of the transaction highlights the primary relationship, and the following ones need only indicate which links to the currently active structure are being worked on.

We come to the position that the major unit of syntactic structure is not the sentence, but the transaction. A transaction may be a sentence, but it may also be a conversation, a monologue, a lecture, part of a conversation, or a single command. It may also include an extra-lingual situational component. Suppose we arrive at the crest of a hill, to see a fine panorama laid out before us. A simple "Beautiful" will suffice to convey the message. The situation is sufficiently similar for speaker and hearer that the speaker may assume that the hearer has the same part of his cognitive network highlighted as he himself does. This view of syntax does not admit the concept of ellipsis. Nothing is missing from a so-called "elliptic" sentence, unless the speaker by design wishes to pose the hearer a problem or set up an ambiguity. Apparent ellipsis merely indicates that the "missing" elements are already available to the hearer and will be correctly fitted by the syntax and sense of what is actually said.

If the major unit of syntax is the transaction, a minor unit must be the relational link, the "agent", "recipient", and so forth. Phrases do not seem to have any great import in this kind of syntactic analysis -- or synthesis, for that matter. A phrase signals that a link is temporarily left open. One of the marvels of language is the enormous parenthetical capability it allows. One can write a sentence, and often, without changing its meaning except for expansions or contractions, can insert or delete words and phrases almost anywhere. Editing a draft paper demonstrates the degree of this flexibility. Inserted words usually have the effect of holding some link already open open a bit longer. Deleting words allows links to be closed more quickly. Moving phrases may reduce the number of links open at any one moment.

The problem of parsing is seen not as how to separate a sentence into the phrases of which it is constructed, but as how to find which links within the cognitive network each word opens and/or closes. Parsing is inherently semantic in nature. Many authors have demonstrated the semantic requirements of syntactic analysis (e.g., Katz and Fodor, 1963), and Schank (1972) in particular has shown how deeply in the cognitive structure these semantic considerations of syntax may be hidden.

Fillmore(1968) took the position that syntax was determined largely by transformations on underlying noun phrases with prepositions. Each preposition was determined by the case of the noun which was related to a verb or to another noun. Sometimes the preposition might be null or deleted by a subsequent syntactic transformation rule. Fillmore's cases correspond to the relationships within the cognitive structure. It does not seem necessary, however, to suggest that there is an underlying prepositional phrase for each relationship. It seems more reasonable to suggest that the content words are themselves primary, and that prepositions are used only when linkages are otherwise obscure. Consider "John gave Jane the book". If one accepts the notion that the programme for "gave" asserts that the links are filled in the order "agent", "recipient", "object", the sentence needs no prepositions. "John gave the book to Jane" requires the preposition "to", since according to the word order, "book" would be the recipient, as in "John gave the book a new cover". Only because "to" signals the recipient of "gave" does the book get relegated to the position of object. "John gave it her" and "John gave her it" are, however, both correct, presumably because of the semantic consideration that one rarely gives an animate being to an inanimate object; although "John gave it her" could substitute for "John sacrificed her to the idol". If this kind of ambiguity were possible in a particular context, the preposition would be required. We see, therefore, that the context can determine the surface syntactic structure of an utterance, and that in this simple structure there are many opportunities for syntactic variation.

One of the major jobs of syntax is to highlight the proper concepts. The first content word of a new syntactic group (transaction, sentence, or phrase) provides the topic for that group. This may well be universally true. It is clearly seen in the topic-comment structure of Chinese (Chao, 1968). English used the passive construction for the same effect, in case what is to be highlighted is the conceptual object of the action: "The book was given to Jane by John.", or the recipient: "Jane was given the book by John". Notice the order of the relational links and the use of pronouns in these forms. When the object comes first, both others need prepositions, but when the recipient precedes the object, only the "misplaced" agent needs its preposition. A similar effect can be made by repeating the topic at the end of the phrase, as is frequently done in French, as well as Chinese. The topic thus highlighted is ordinarily present in the pre-existing concept structure, and the highlighting then serves to link the new information into its correct place in the network.

In English, highlighting is often performed by making the topic into the first noun, rather than the first content word. Consider the difference between "A red chair..." and " A chair-shaped redness...", or among "John left and Mary was sad" and "John's leaving made Mary sad" and "Mary was sad because John left". The actions and events are the same, but the topic is not. Later items in the transaction will more probably enlarge on the topic than on the other items in the utterance. The topic is the highlighted concept. A good part of syntax is needed to keep straight the linkages whose "natural" order is violated by the need to highlight particular topics. The rules whereby this is done are very different across languages, and these differences give learners of ne. languages major difficulty.

Negative and uncertain concepts.

A facet both of perception and of language is negation. To negate something is not the same as to ignore it. When a state of the world is negated, anything else may be true, except for the specific potential state which is negated. Perhaps this is the essential fact about negation: the things which are possibly true about the situation are ordinarily more varied than the thing which is not true. If the reverse holds in a particular situation, it is usually more economical both perceptually and linguistically to assert what is or may be true. The smaller and more precise of the two distributions (false states and true states) will usually be the one asserted. Negation therefore should be indicated in the cognitive network by the existence of a structure which signifies a positive state of some kind; this structure will be linked to the rest of the network by a special relational link which negates the import of the structure.

Suppose we have "John did not give Jane the book". This is a very ambiguous statement in English, since we do not know exactly what is being denied. Clearly, the entire statement "John gave Jane the book" is not true, but it may be true that Jane was given a book, that John gave someone a book, that John gave Jane something, the Jane stole the book from John, or that there was no transaction involving any of the participants and giving. The cognitive structures related to these different possibilities differ.

In the following figures, the sign % indicates that the concept marked is asserted not to exist within the pattern. In a pattern matching to check whether the structure exists elsewhere in the cognitive network, the match will succeed if another structure is found to be the same except for the concepts marked %, and will fail if the structure contains the concept. For example, < %John <--agent-- went> will match <Bob<--agent-- went>. Fig. 13 indicates some of the different possibilities for "John did not give Jane the book".

%John <--- gave ---> Jane ↓ book It was not John who gave Jane the book.	John <--- gave ---> Jane ↓ %book It was not a book that John gave Jane.
John <--- %gave ---> Jane ↓ book John did not <u>give</u> Jane the book.	John <--- gave ---> %Jane ↓ book It was not Jane to whom John gave the book.

Fig. 13

The null transaction, where there was no giving or other relationship between John and Jane, cannot easily be signalled in this way. Each of the structures in Fig. 13 suggests that something positive happened, but it did not involve the negated concept. Substitution of some other concept in the same frame would lead to a true statement. In English, the null transaction can be stated explicitly as "John gave Jane nothing". This way of putting it states that there is no concept that can correctly be fitted into the frame, since all possible objects will match "thing" in Fig. 14. Similar denials of the whole transaction may be made by asserting that there is no match to any one of the concept elements:

```
John <--- gave ---> Jane     %person <--- gave --->Jane     John <--- %act --- >Jane
           |                               |                          |
           v                               v                          v
        %thing                            book                       book
```

Fig. 14

In English, the structures of Fig. 14 correspond respectively to "John gave Jane nothing", "Nobody gave Jane a book", and "John did nothing with Jane and the book". The last assertion is considerably stronger than the others, denying any kind of triple relationship among John, Jane and the book.

The conceptual negations probably are language independent. They should be, if the basic structures of the network, like "agent" "recipient" are themselves independent of language, as we theorize. The syntactic realizations of the structures are obviously dependent on language. English seems to prefer to negate the verb. Indeed, it is not clear that there is a natural method in English of negating anything other than verbs and modifiers. If we want to say "Not-John gave Jane the book", we must say positively, "It was not John who gave Jane the book" negating "was". The corresponding surface cognitive structure is as in Fig. 15.

```
John <---- %isa -----<        ><-------- gave ------->Jane
                       |                   |
                       | isa               |
                       v                   v
                     person               book
```

Fig. 15

It would not be unreasonable to suppose that cognitive structure handles negatives this way, but it is hard to see how to test such an assertion. Do other languages show the same restriction of negation preferentially to verbs and modifiers that English shows?

Uncertainties in structures may be handled the same way as negatives. One can imagine that the various phrases "it was probably John", "it might have been John", "was it John?" and the like could each substitute a marker for the % indicator in the forgoing. The idea again is that the positive structure is linked to the rest of the network by some quantifier, truth assertion, or question indicator.

Language and the cognitive network; summary.

Initially, the cognitive network grows from recurrent patterns in the environment. At some stage, these patterns have formed concepts of sufficient generality that they represent classes of objects for which adults have names. The names will be heard by the child on numerous occasions when the concepts are highlighted, and will thus come to be associated with the concepts. When this has happened with enough names and concepts, the relationship in which a sound is linked with a concept will come to form a concept in its own right, that of "label". Other words will then be rapidly learned as labels for new concepts, and the child should seek labels for his unlabelled concepts.

By a similar, probably simultaneous, but logically unrelated process, the child may learn that when he makes certain sounds, particular concepts come to be highlighted. "Mama" brings mother, for example. At some stage, this relationship also will come to be a concept "output label", and the child will attempt to place labels onto his concepts. At around this stage, the matching of input and output labels should begin. Initially, the output labels may be idiosyncratic, but soon sould become similar to the input labels because the adults will ordinarily respond only to words which resemble their own. The child learns adult speech more than the adult learns the child's, although the latter process does occur to some extent. When the two label concepts have been stabilized, the child's output label serves also as an input label which sounds something like the label he discovered in adult speech. As he varies the exact sound of his output labels deliberately, or because of inadequate articulatory control, sometimes his words will match the adult sounds better than at other times. On these occasions, the linkage between concept and label sould be reinforced more than when his word is aberrant. Hence the child's speech patterns come to resemble those of his community.

Syntax may be learned by a very similar technique. Relationships among concepts come to be concepts in their own right. "Above", "inside", and so forth, are perceptual relationships which become concepts. Later, more complex relational concepts probably are made, like "agent" and "object". These are both more common and more abstract than the object concepts they relate. Matching relationships occur in adult speech. Each relationship in the cognitive network potentially has a syntactic indicator, which can be learned in the same way as the label is learned.

When the network is mature, the syntactic structures may be used freely to indicate which links in the cognitive network are to be filled or opened by each word. Most of this linking job is handled by the meanings of the words in the total context of the utterances, but syntax provides a fail-safe backup system in case the meanings alone do not ensure correct matching. Large-scale syntactic structures, like phrases, may be used to indicate that one or more links are held open while a related structure is being built. The appropriate unit of syntactic analysis is not the sentence, but the transaction, since the transaction defines the context, and links remain open across sentences within a transaction. In written language, the opportunity does not arise for feedback to ensure that linkages are correctly filled. Hence, written language must be structured differently from spoken language. Each small segment must be expected to define its topic more carefully than in spoken speech, and must be more careful to close links which in speech might be left open for the other party to query if he wished.

As an overview, perception might be considered to be the foundation on which the cognitive network is built. Perception can activate concepts from below, so to speak. The further one is from the perceptual foundation, the more abstract the concept, the more different the conceptual structures of different people. Language, on the other hand, can access any concept at any level of the network, provided that there is some concensus among people as to the meaning of the concept. Structures laboriously built by perception may be readily transferred to someone else by the use of language. Language, to follow the metaphor, gives a side-view X-ray of the structure built on perception. In the following, we shall examine the role of different languages within the same network.

Output labels are more difficult to deal with. We discussed output label acquisition in several ways, including Skinnerian shaping of autonomous speech, and the use of feedback between the child's output and his already acquired input labels. If two different labels for a concept must be acquired, the Skinnerian approach demands that each be separately shaped. The label in one language must be evoked by the concept; but so must the label in the other language. Some other factors must determine which of the competing responses actually occurs. These other factors presumably include social and verbal context. Labels in one language are more likely to be reinforced in the presence of one adult than another, or when other particular words have been used. "La tête" is more likely to be heard by the child when "le bébé" has been used than when "the baby" has occurred. Accordingly, the child's internal feedback mechanism should accept "head" more readily as a correct output label after someone else has said "the baby" than if the context was French.

These influences are likely to be weak at first. We are talking about the period in which the child is first learning labels. Conditional acquisition should be considerably more difficult than simple acquisition. We may expect, then, that not only should the learning of vocabulary be delayed in a child reared bilingually, but also the differentiation of the vocabulary according to language should follow by quite a while the acquisition of the words. In a truly balanced bilingual, if such exists, the first words should be randomly chosen from either language.

Most bilingual children are not balanced. They may be equally capable in either of their languages, but ordinarily they may associate one language with one set of their adults, and the other with another set. Mother may speak mainly French, father mainly English, and so forth. If there is a pure environmental separation of this kind, it should be conditioned as part of the labelling pattern. Instead of the pattern of Fig. 9, we might expect the labelling pattern to look more like Fig. 16.

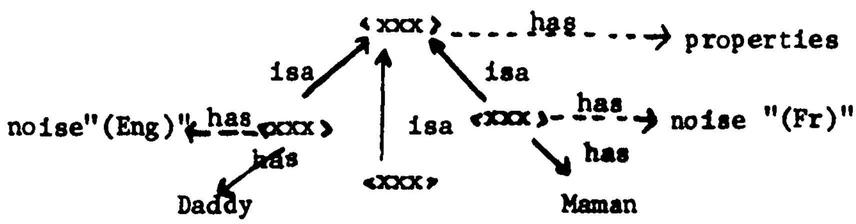

Fig. 16

Fig. 16 explicitly represents the input labelling process which might occur if Daddy spoke English and Maman spoke French, but because of the self-corrective feedback from output to input labels, it also represents a mechanism which may help condition the use of output labels according to the person being spoken to. There may, in the same social context, be a direct reinforcement of the separation of languages if Mummy does not respond properly to English labels and Papa fails to respond to French.

A further mechanism for label segregation, but one which does not differentiate between the two languages, is unbalance in the labelling for each individual concept. A concept may be labelled in one language but not in the other. It can only be spoken in the one language. If the child has reached the syntactic stage of speech, rather than the single word stage of which we are talking, his attempt to use the concept while speaking in the other language may result in what an adult sees as an "intrusion". A word in one language is inserted into speech in the other language. But this is not intrusion from the child's viewpoint. There is no label pair for that concept, and thus no choice. At this early stage, words have not been differentiated into languages. There may be "Mummy words" and "Daddy words", but if a concept is labelled only once, the label may well not be differentially conditioned to either parent. Particularly if the conversation is with neither parent, there may be no reason for the child to even recognise that there is any dicrepancy in his speech. His input label conditions will not differentiate between the languages in the absence of the adults who provide the conditions.

The discussion has proceeded in terms only of object concepts. The same kinds of argument may be made in respect of the relational constructs which predominate in the second stage of the child's single-word development (Bloom, 1973). Indeed, the argument may be stronger in respect to concepts like "up", "more", etc., since they occur more often than do the particular objects labelled in the first stage. Furthermore, it is by no means clear that the object concepts have stabilized in the adult form during the early single-word stage. Labels such as Guillaum's "nēnin" (quoted by Bloom, 1973), which referred at first to the breast, then to a wide variety of things like a red button, a pointed elbow, and a picture of mother, seem to label concepts at a much more concrete level than do adult labels. They label each of the features of an experience with the same word, and do not differentiate among the various features. Words learned at this stage seem to have a tendency to drop out of use during the second stage, perhaps as the concepts become more refined, perhaps as the relational concepts become better and more stably developed. By the time they return, and the vocabulary take-off commences, the labels refer to more abstract concepts, which more or less agree with the adults idea of concrete objects.

In the bilingual infant, the same kind of shifting representation should occur. The very early bilingually labelled object concepts should come and go, while the more frequently useful syntactic relations (perceptually syntactic, that is, like "up") should appear and be stabilized in both languages. Since these are the very words which tend to occur in the first syntactic utterances, they should have a considerable influence on the differentiation of the child's speech into two separate languages.

Language segregation in the bilingual child

The development of syntax should herald the segregation of the child's language into two separate languages. At the syntactic stage, the verbal environment of words begins to become important to speech recognition and production. The perceptually syntactic concepts begin to appear in relation to object words: "more cookie" "up chair", and so forth. These relational words have been used for some time by the child, but only as isolated words, not in a syntactic way. When they become associated with object concepts, the question arises as to how they become associated with labels of the proper language. There seems no initial reason why they should preferentially associate with their "proper" language than with the other, except that the social context is biased the same way for both. "More" and "cookie" are both "Daddy words", whereas "encore" and "(Fr) biscuit" are "Mummy words".

Once words come to be paired syntactically and preferentially by language, they come to have mutual associations. Furthermore, the syntactic programmes which develop within the cognitive network as properties of the perceptual syntactic relationships will have their own biases as to social context, and hence will be language-biased. The mechanism whereby this happens should be the same as for labels. In contrast to labels, however, syntactic programmes involve operations with words, and the words themselves are language-biased. We still cannot say that the language is a property of the child's words, but both the labels and the syntactic relationships are simultaneously subject to common social contexts, and their biases should reinforce one another.

After a certain proficiency with syntax has been achieved, the syntactic patterns which have been learned will show common biases. For example, there will be a whole class of object words which are sometimes heard after "green", and a separate class heard before "vert" in the same perceptual context. All of the first class are "Daddy words" and all of the second class "Mummy words". The self-corrective feedback from output patterns to input patterns should slowly permit the child to discover that putting a Daddy word before either "vert" or "green" does not sound right. When the individual syntactic patterns begin to be correct more often than not, then and only then should the child begin to discover that there are two separate languages and pattern his speech according to them.

The argument here is that there is no "language switch", as such. Rather, there are a large number of individual relationships among words, which together form the syntax of two languages. The words which fit a syntactic relationship of one language often fail to fit the other, and hence utterances are usually confined to a single language. Failures to keep to a single language will often occur because the appropriate label i- not known in the intended language, or because the other label will fit into the same syntactic structure. Similar "syntactic" types of relationship may be expected at the other levels of the concept hierarchy, resulting in phonological, verbal, syntactic and conceptual consistency of language by the time that the conceptual levels have become fully stabilized.

Children usually begin to use syntax of a kind by their second birthday. Nevertheless, Michael, the three-year-old studied by Swain and Wesche (1973) made intrusion errors (as oppored to immediate retranslations) in both directions between his two languages. These errors declined in frequency, and were more often spontaneously corrected during the second period of the study (after age 3.5). The overt occurrences of language mixing of all kinds were few, amounting to no more than 4% of all multiword utterances recorded, which suggests that the separation of languages had been well accomplished before the start of the study. Nevertheless, the fact that corrections did not become frequent until later suggests that the syntactic self-corrective mechanism was not the primary influence for segregation in the early stages. Words were recognised as belonging to one language or the other only after they were almost always correctly used.

Michael was not a balanced bilingual. His English was less well developed than his French. Swain and Wesche comment that his English had the "telegraphese" character of early syntactic language, whereas his French was more mature. The substitution errors conformed to this pattern, only a few of the French words in English sentences being nouns and noun phrases, whereas most of the English substitutions in French sentences were of these classes. Michael's tendency was to use French syntactic structures and markers, regardless of the language of the content words. According to the analysis above, he would not be expected to distinguish between his languages at this stage, since only the French syntax had begun to be stabilized. He apparently began to correct and to hesitate while looking for words of the correct language only after his English syntax began to grow, as the cognitive network viewpoint would suggest should happen.

Second language learning

Very few bilinguals, even among those who, like Michael, grow from infancy in a bilingual environment, are balanced in having equal ability in both languages at all levels of development. It is probably as true to say that all bilingualism involves some degree of second language learning as it is to say that all language learning can be considered as a unity. Different dialects and jargons or even styles within a single "language" require to be learned in addition to the basic language, and there is probably no real difference between such learning and the learning of a totally new language.

Whether two languages are learned together or separately, the final requirements are the same. Words and syntactic programmes of two separate languages must be linked together properly and must also be linked to non-linguistic conceptual structures. There is a difference, however, between those cases in which the structures of the second language are being built while the structures of the first are still fluid and those cases where the structures of the first language are rigidly stabilized when the second is being learned. In the former case, the initial learning is difficult, because the separation of the structures of the two languages must be conditioned to environmental factors outside the language, such as the presence of Mummy. In the latter case, the first language is more readily learned, but the second is much more difficult than when both languages are learned together. Not only must new labels be attached to concepts which already have labels, but also syntactic programmes must be built where perfectly satisfactory means of expressing the conceptual relationships already exist. Both the experience of Michael, noted above, and the findings of Dumas, Seliker and Swain (1973) on the prevalence of English syntactic structures in children learning French through an immersion course attest to the difficulties of adding a new syntax to a pre-existing one. French people passing through Toronto have commented that the French heard on radio station CJBC (the CBC French language station in Toronto) is structurally English, by which they probably mean that is is more influenced by English than is the French spoken in France.

Why should syntactic patterns be difficult to reduplicate? One possible reason is that syntax is not required for comprehension. Telegraphic speech can be understood, provided that the content words are arrayed in some reasonable way. Syntax, as we discussed in Part II of this paper, serves mainly to keep track of open links in the cognitive structure, and to highlight concepts at the speaker's pleasure. If syntax is not required for understanding, the second-language learner can make his needs known by the use of elementary syntax, which may be based partly on simplified rules for the target language and partly on well-known rules for the mother tongue. Once he reaches this stage, motivational factors come strongly into play. Little is learned in the absence of attention and motivation. If syntax is needed only for elegance, only those motivated by elegance will learn it elegantly, once they can communicate.

With total immersion in the milieu of the new language, input syntactic structures will, of necessity, grow so that the speech of others may be understood. Self-corrective feedback will slowly tend to make the output syntax conform to these structures, but without deliberate effort, the process will be slow. How many adults know the sound of their own voice? One hears what one wishes to produce, not what one actually produces.

Training for a second language

Different schemes have been used for various purposes in training people in a second language. There are four separate routes for information into and out of the cognitive network; different teaching methods emphasize different combinations of these routes. The classical school technique of learning vocabulary lists and rules of grammar represents the purely verbal input route. Language is used to convey the cognitive structure of the target language as it is known to linguists. By this method, a section of cognitive network, which, like a map, represents the target language, is built in the learner's larger network; but no action programmes are built which would enable him to understand or to produce utterances in the new language. Those he does produce are done more like a mathematical exercise than like meaningful communication. Existing programmes are activated by complicated routines working on structured knowledge, whereas the desired state is for new programmes to be activated according to the needs of communication.

One can learn a new language by the classical technique, but it needs practice. The motor and language output routes must be brought into play. The knowledge learned by rote can be used to produce and to decode utterances in the target language (probably only in writing, speech being too fast for the complex processes required). It can also be used in conjunction with external feedback to provide a weak kind of self-corrective feedback which might help the learner build the correct active programme structures. In the end, a motivated learner could probably learn a new language quite well by this method, at least for reading and writing.

The "Voix et Images" method might be considered to be the extreme opposite of the classical method. In this technique, only the perceptual and motor routes are used. Language is not used as a route for conveying information, but the learner's attention is drawn to facets of the target language which he must learn by example, and by discovering perceptually what the new noises refer to. It is an attempt to reproduce the learning situation faced by the monolingually raised infant. But adults are not infants, and the "Voix et Images" technique seems to miss the possibility available to adults of using language as a vehicle for conveying structure from one cognitive network to another.

Immersion training is an extension of the "Voix et Images" method, but lasts for longer periods. Teachers refrain from speaking anything other than the target language, and the learners are expected to do the same. They learn from the linguistic environment, ad lib, rather than in a restrained and directed way as they do with the "Voix et Images" method. In terms of the cognitive network, sustained immersion in a monolingual environment should be adequate to permit anyone to obtain at least an ability to communicate in the new language. Problems arise, however, when structures in the old language seem to correspond with those in the new, but in fact represent something different. Such things must be found by the slow "perceptual" techniques. Errors in the new language may well persist for a very long time under such circumstances. Immersion in a class of learners who are all at the same stage should cause problems. Most of the interactions are with the class members, whose vocabulary and structures are not those of the target language. A phenomenon akin to "twin language" should be expected, in which the language learned is neither the old language nor the intended target.

In considering second-language training from the viewpoint of the cognitive network, one must think first and foremost about the stabilization of levels. It seems reasonable that second-language learning should commence at the most accessible place in the cognitive network. In the pre-lingual infant, this is probably the phonological level and the level of concepts conveyed by language. To train a child, you expose him to the sounds and words of the target language. To train an adult, you tell him what you are going to try to get him to learn.

Children cannot build grand structural edifices of subtle syntax and abstract vocabulary. They do not have the foundations of simple syntax and concrete vocabulary. Similarly, adults only with difficulty can restructure the simple sounds they emit, and learn the fundamental syntactic structures anew. While it is probably useless to practise a child who has only a rudimentary idea of conditional concepts in the use of the subjunctive, it is probably necessary to drill the adult in hearing and producing new phonemic distinctions and in using new syntactic structures for the basic perceptual syntactic relationships. It seems to follow that most children should at least be exposed to the sounds of many languages, so that they may have the phonemic structures available in adulthood, when they may want to learn a new language.

Vocabulary must be considered separately from phonology and syntax. The concept structures of all people are open for the construction of new vocabulary. The main problem arises with the fact that words in different languages do not mean the same things, either denotatively or connotatively.

In Russian, the same word "krasnyi" is used for "red" and "beautiful", and in its "red" meaning has much of the connotation we reserve for "gold" -- goodness, purity, and so forth. The Red Army seems very different to a Russian.

It is often said that translation is impossible. But vocabulary in a second language is often learned by translation, even if the training technique attempts to avoid it. The translation occurs either because a new word is linked to a pre-existing word, as happens when the translation is given overtly as the meaning of the new word, or because the concept to which the new word is linked as a second label is a pre-existing concept labelled by an old word. Only by prolonged observation or by direct instruction can the learner discover that the concepts in the two languages should be different. Even then, it seems likely that the old-language concept will survive in the new language as a strong associate, which will be activated along with all its connotations when the new-language concept is evoked. We shall return to this point in considering linguistic relativity and the Whorfian hypothesis.

Third-language learning

Third and subsequent languages are often supposed to be learned more readily then is the second language. The cognitive network approach suggests why this might be so. It should be true for distinct but related reasons.

In the earlier passages, we have frequently referred to the duplication in the new language of syntactic structure for a given perceptual syntactic relationship. The syntax appropriate to the new language is somehow grafted onto a perceptual structure which already is linked to a syntax programme. The same applies to double labelling. One problem in this process has been somewhat glossed over. The pattern in which a single monolingual label or programme is linked to a concept is not the same as the pattern whereby one of a polyglot set of programmes or labels is linked to its concept.

The difference in patterns between monolingual linkages and multilingual linkages between concepts and their labels or syntactic programmes is that in the monolingual case the link is direct, whereas in the multilingual case, each link is contingent on some external condition. This conditional linkage was brought out in connection with the "Mummy words" and "Daddy words" of the bilingual infant. Making such linkages is harder than making simple unconditional linkages. Harder yet is the construction of a completely new pattern which must supersede an old one. When a monolingual learns a second language the general pattern of <conditional linkage> must replace the pattern of <direct linkage> When a bilingual learns a third language, there is no new pattern type to be learned. This should materially assist the learning of extra languages. However, if the new language differentiates structures in a way that the other two do not, then it again requires the learner to build new structures in place of old ones, and the difficulty should arise anew. It is always easier to add structure to the cognitive network than to reconstruct. Indeed, it is quite possible that the original structure is never lost, but is merely overlaid by the new.

Another way in which learning a third language should be easier than learning a second is in the building materials at hand. Consider the phoneme structure as an example. In either of the first two languages, each phoneme covers a certain region of phonetic possibilities. These regions do not match exactly in the two languages. If the learner has properly learned the phonemic structure of his first two languages, he should have available at least the features which enter into both patterns of phonemes, as input to phonemic structures to be built for subsequent languages. In other words, he can hear better than can a monolingual. The same applies at higher levels of the cognitive network. Differentiation of features across the original languages provides a richer vocabulary of features on which to build a new language. Each subsequent language should be easier to learn. Difficulties will still arise with new languages which make use of features ignored by all the original languages, both because the new features will not be in the available vocabulary, and because the necessary conditional linkages will not have been constructed. The learner will both have to learn to hear the feature (Chinese tones, for example), and to set up a conditional pattern whereby he hears them only in context of the new language. He does not want to spoil his understanding of other languages by taking account of a nonsense feature.

Third languages should be easier to learn than second, but shifts into new language families should be nearly as difficult as learning a second language in the same family as the first.

Bilingualism and linguistic relativity

Does a person's language affect the way he thinks? The strong positive statement of this hypothesis is known as the Whorfian Hypothesis (Whorf, 1956). On the other side, most experiments aimed at the question have provided negative results. The cognitive network idea seems to suggest that linguistic relativity (the Whorfian Hypothesis) should apply, but not to concepts generated largely by perception. The more abstract the concept, the more likely it is that the language embedded in the network is implicated in the construction of the concept. Someone has used it in conversation, for example.

The concepts in the cognitive network are all generated by the discovery of relationships among other concepts. At the most primitive level, the basic concepts may be "line", "edge", "colour", and such elements. At a little higher level, concepts should be based on these primitives, e.g. "ball", "face", "sky", and should include relationships, e.g. "up", "in", etcetera. Slightly higher-level concepts involve relations among the lower concepts, e.g. "give", "move", "big". And so forth. After many levels of relationship, each level building on the relations already secured, the concepts become really abstract, e.g. "phenomenalism", "trans-substantiation". Along another line of abstraction, mathematics, concepts such as addition and multiplication give rise to functionals, mapping, groups, and so forth. The language of mathematics is particularly framed to aid abstraction. Relationships are carefully built on relationships, to a degree of abstraction and precision unmatched by natural language.

The cognitive network may be visualized in one way as a vast structure of linked nodes, resting on a foundation of sensory data. The higher in the structure, the wider its reach, and the more idiosyncratic the individual links. At a very high level of abstraction, probably no two people maintain the same concepts. Each level up the structure provides some opportunity for a slightly different arrangement of concepts to connect the concepts of the layer below.

Language provides a different skeleton for this structure. Any concept, from any level of the structure, may be labelled or provided with an action programme. Once the concept is labelled, its pattern can be transferred to another listener. Provided that the concepts in the pattern are common between the speaker and the listener, the transfer of the new pattern should be reasonably successful. Such agreement can occur in the first instance only if the basic elements of the new pattern are low in the hierarchy, where perception provides the main base of stability. When the labels for these patterns have been well established, then more complex ideas can be transmitted. In natural language, however, there is always some uncertainty as to the range of properties implied by a concept. Every concept is linked to every other one in some way, and connotative meaning is passed along with the more tightly clustered denotative patterns. Hence, it takes years of study to get an agreed meaning for a conceptual label such as "trans-substantiation", and people spend their lives on such matters.

In this context, it would be surprising if different communities did not hold somewhat different meanings for common words. (Consider the word "management", from a shop steward's and a company director's point of view.) It would be still more surprising if people brought up to use different languages did not have a consistently different set of concepts at reasonably abstract levels to cover the same range of meanings.

Language can stabilize the cognitive network from person to person at some level of abstraction above the level where pure perception could yield inter-personal consistency, but it cannot stabilize it all the way. Accordingly, we should expect to find three domains of the effect of language on thought. At the lowest level, perception rules, and all people should have the same basic conceptual structure. At an intermediate level, people in basically the same kind of environment who talk to one another should have the same sort of structure. Both environmental variation and a lack of communication should lead to real differences in the way different groups conceive their world. At the highest level of abstraction, each one of us has his own ideas, and no way of knowing, without skilled effort, whether or not these concepts are shared by anyone else.

Linguistic relativity is thus a consequence of the fact that people having different languages usually do not speak to each other much, and to the fact that they usually have different environments, which tends to stabilize even the lower-level concepts somewhat differently. Mathematicians in a given field should be able to work together at much higher levels of abstraction than can people who attempt to use natural language. The ideas of physicists are almost entirely communicated and directed by the form of mathematics available to them. These concepts have profound influence on the course of the world. The mathematics of tensor calculus in Einstein's mind was responsible for the Nuclear Age. Different forms of mathematical language have been responsible for almost all the really important social changes in the last hundred years. It seems that Whorf was just too timid in asserting his hypothesis. Not only does the structure of language affect the way we think, it controls the whole content of our next-to-highest-level thoughts. It cannot control the really creative thoughts, because these require mental manipulation of concepts into new patterns by one person. But the communication of the results of creation is impossible unless the concepts on which the creation was based are already in the language.

The implications of this analysis for the thought of the bilingual are clear. If the bilingual is an adult, whose conceptual patterns were developed before his second language, then most of his conceptual substructure will be stabilized by his first language, and it is as difficult to restructure these levels of the network as it is to rebuild the fourth floor of a forty-story building. A few new concepts may stick, but they will be built into the old network. If, however, the new language is learned in a new social context, a whole new body of cognitive structure will be developed with it, and may well be responsible for the apparent language-dependent personality changes observed by Ervin (1964) and Ervin-Tripp (1968). The situation is quite different for a bilingual infant. He is building his conceptual network at the same time as he is learning his languages. Both languages will stabilize the network, and middle-level concepts from each language structure will be employed. Whether this means that he will have a structure that is like that of neither of the parent structures, or one which employs most of the elements of both, is moot.

Certainly, his structure should be more flexible than either of the parent structures, since the concepts appropriate to one language will not quite properly relate to concepts of the other, leaving room for change. Either he will have a richer conceptual substructure than a monolingual, or he will have a more flexible one. Both possibilities seem to suggest that a person who becomes bilingual as an infant should have a greater opportunity to be truly creative than one whose bilingualism was acquired later.

GENERAL SUMMARY

It is as well to remember that the title of this paper is "Speculations on.....". These are indeed speculations, but they are not fantasies. Consideration of the growth and structure of the cognitive network provides a coherent viewpoint on the nature of language as a whole, and on the problems of bilingualism in particular. Even if it is correct in its major outline, it is bound to be false in detail, and grossly oversimplified where correct. Nevertheless, it seems to be helpful on a partical level.

The growth of the cognitive network is based on three main principles: (i) if a pattern recurs often enough, it will be segregated and stored as a concept, to be reactivated when the same pattern of properties or events recurs; (ii) concepts previously constructed may enter as elements of patterns of new concepts; (iii) as the network grows, the concepts derived earlier become stabilized. New concepts may be added, but it is hard to alter old ones.

The three principles of network growth are derived from ideas about pattern recognition in general, and in various combinations have been used in the construction of automatic pattern recognition devices, as well as for theoretical constructs. A novel feature in the network as used here is the inclusion of transformational operations as concepts within the network, and their identification with both perceptual and linguistic syntactic structures.

The cognitive network contains all our knowledge of the world, particular and general, data and programmes, our information and our abilities. In particular, the principles which govern the growth of the perceptual part of the network also govern the part which deals with language. Labels relate words to concepts, and syntactic programmes apply similarly to relationships among concepts. Most syntactic relationships deal with only two concepts at a time, but some, mostly indicated in English by verbs, deal with three. Syntax is taken to be the primary means whereby confusions about how concepts link together are resolved. Syntax also serves to ensure that the intended concepts are highlighted and keyed into the existing cognitive network. As a general statement, syntax serves to aid in putting new information into its proper place in the cognitive network, and is an accessory part of language, not its primary attribute.

According to the cognitive network viewpoint, language should not be analysed sentence by sentence. Rather, an entire transaction should be analysed, and its syntax related to the currently highlighted and the currently open aspects of the conversation. Written language, however, has a different formalism. Because the reader is not present, and is unable to raise questions when a point is obscure, written language has evolved so that each sentence restates its connection with the rest of the cognitive network, even though this restatement may be redundant and unnecessary in conversation. As a corollary, written English sentences almost always contain a verb, because a verb is the only relational concept that can link two other concepts in a self-contained piece of information.

In respect of bilingualism, the problems of learning two languages from infancy can be differentiated clearly from those faced by a monolingual learning a second language. If an infant is exposed to two languages more or less equally, then his cognitive network will build labels and syntactic programmes corresponding to both languages. The languages will initially be undifferentiated or at best will be conditioned to the presence of specific other people in the environment. Only when the child has stabilized a variety of syntactic relations will a language be segregated by virtue of the verbal context. Even at this stage, words of the less stable language should be found in syntactic frames of the primary language. The concept of the "language switch" has no meaning in the context of the cognitive network. Each language is spoken coherently only because all its components are more strongly related to other aspects of the same language than to those of the other language. Environmental conditioning serves to aid segregation even after syntactic segregation has been achieved.

The stabilization of concept levels during the growth of the network suggests that adults should have a very different problem. Whereas the bilingual child is learning the phonology of both languages at a time when his phonemic repertoire is fluid, and hence can learn both phonological systems at once, the adult is already provided with a stable phonological system. He must painfully restructure this low concept level in his cognitive network, both for input and output. The same applies to the fundamental syntactic structures.

Being low level structures, they tend to be resistant to change. Adult second-language training should concentrate on breaking down old structures which conflict with the requirements of the new language, and to this end should use all the tools at hand. These tools include the adults' existing ability in his native tongue. Things can be explained to him. They include practice in unfamiliar phonetic and syntactic structures. Language labs and immersion courses do this. The problem is completely one of adding and substituting structure in the cognitive network at normally inaccessible levels of the net.

The cognitive network idea sheds some light on the question of linguistic relativity and on whether language precedes or follows the development of cognitive structure. The network is originally based on perceptual information, manipulated mentally. It must be possible to develop concepts independently of language. But once the ability to communicate by language has been acquired, structure may be imported wholesale from other people's cognitive networks. Naturally, someone must have first constructed the patterns to be transmitted. Also, if the concepts on which the new structure is built do not exist within the old network, then the new structure is not learned. Communication will have failed. Provided that the state of the network is ready to accept it, a new concept can be imported through the use of language. But creative concepts can not be imported. They must be invented. Language thus affects only the middle regions of the network.

The network is stable at the lower levels, and being based closely on sensory data, is common to all people, and to some extent to all related species. At a slightly higher level, language directs the structure of the network, but is based on experience common to all, and should thus be common across all languages. More abstract concepts should be common across members of a linguistic community, but not between communities. Here, linguistic relativity should begin to come into play. The language drives the way people construct their concepts, and the concepts control the development of new language. At a still higher level of abstraction, there is no language, and everybody creates his own world view. Mathematics is a powerful language created to handle abstraction, and through the use of language the events that shake civilization have been born. Only the concepts of the physical world seen through the eyes of mathematicians have made the nuclear age, the transistor, and biodegradable plastics.

Whorf understated his case for linguistic relativity, which may be the most important fact of the world today. Mathematicians do not think as common men, when their creations are based in concepts spoken only in mathematics.

Footnotes

1. This is D.C.I.E.M. Research Paper Number 74-RP-1013.

2. I thank Dr. I.K. Taylor for assistance with this paper, and Dr. M. Swain and G. Dumas for their comments on an earlier draft. None of these can be blamed for the many known omissions which occur, particularly in the section on bilingualism. I felt that the readership of these "Working Papers" would probably be more familiar than I with the research work on bilingualism and second language learning, and therefore chose to concentrate on the implications of the cognitive network viewpoint rather than on corroboration or refutation of these implications. It would take another paper at least as long as this one to deal appropriately with the evidence.

3. These applications will be discussed in the forcoming book, Norman, D.A. and Rumelhart, D.E. Exploration in Cognition, San Fransisco: W.H. Freeman, 1975.

4. I have argued informally that the number of different kinds of link can be reduced to one, a link which activates or inhibits the concept to which it points when the concept from which it springs is activated. For the purposes of the present paper, though, it is sufficient to reduce the number of different links to two. The further reduction to a single link is rather more complex.

References

Abelson, R.P. and Reich, C.M. Implication modules: a method for extracting meaning from input sentences. Proceedings of the International Joint Conference on Artificial Intelligence. Washington, D.C., 1969, 641-647 (Quoted by Rumelhart and Norman, 1973).

Blakemore, C. and Cooper, G.F. Development of the brain depends on the visual environment. Nature, 1970, 228, 477-478.

Bloom, L. One Word at a Time. The Hague: Mouton, 1973.

Burling, R. Language development of a Garo and English-speaking child. Word, 1959, 15, 45-68. Reprinted in Ferguson, C.A. and Slobin, D.I. (Eds.) Studies of Child Language Development, New York: Holt, Rinehart and Winston, 1973.

Chao, Y.R. A Grammar of Spoken Chinese. Berkeley: University of California Press, 1968.

Dumas, G., Selinker, L. and Swain, M. L'apprentissage du français langue seconde en classe d'immersion dans un milieu torontois. Working Papers on Bilingualism, 1973, 1, 66-82.

Ervin, S.M. Language and TAT content in bilinguals. Journal of Abnormal and Social Psychology., 1964, 68, 500-507

Ervin-Tripp, S.M. An analysis of the interaction of language, topic and listener. In J.A. Fishman (Ed.) Readings in the Sociology of Language. The Hague: Mouton, 1968.

Fillmore, C.J. The case for case. In E. Bach and R.T. Harms (Eds.) Universals in Linguistic Theory, New York: Holt, Tinehart and Winston, 1968.

Katz, J.J. and Fodor, J.A. The structure of a semantic theory. Language, 1963, 39, 170-210.

Leopold, W.F. Speech Development of a Bilingual Child: A Linguist 's Record. Evanston, Ill.: Northwestern University Press, Vol. III, 1949.

Liberman, A.M. Some results of research on speech perception. Journal of the Acoustical Society of America, 1957, 29, 117-123.

Lindsay, P.H. and Norman, D.A. Human Information Processing. New York: Academic Press, 1972.

Neisser, U. Cognitive Psychology. New York: Appleton-Century-Crofts, 1967.

Pettigrew, J.D. and Freeman, R.D. Visual experience without lines: effect on developing cortical neurons. Science, 1973, 182, 599-601.

Rumelhart, D.E., Lindsay, P.H. and Norman, D.A. A process model for long-term memory. In E. Tulving and W. Donaldson (Eds.) Organization of Memory. New York: Academic Press, 1972.

Rumelhart, D.E. and Norman, D.A. Active semantic networks as a model of human memory. Center for Human Information Processing, University of California, San Diego, Report CHIP 33, 1973.

Schank, R.C. Conceptual dependency: a theory of natural language understanding. Cognitive Psychology, 1972, 3, 552-631.

Stevens, W.J. Verb Types in modern English. Language Sciences, 1973, 26, 29-31.

Swain, M. and Wesche, M. Linguistic interaction: case study of a bilingual child. Working Papers on Bilingualism, 1973, 1, 10-34.

Taylor, J.G. The Behavioural Basis of Perception. New Haven: Yale University Press, 1962.

Taylor, M.M. The problem of stimulus structure in the behavioural theory of perception. South African Journal of Psychology, 1973, 3, 23-45.

Whorf, B.L. Language, Thought and Reality. New York: Wiley, 1956.