

STANFORD ARTIFICIAL INTELLIGENCE PROJECT

MEMO AIM-162

STAN-CS-265-72

PRIMITIVE CONCEPTS UNDERLYING VERBS OF THOUGHT

BY

ROGER C. SCHANK
NEIL GOLDMAN
CHARLES J. RIEGER
CHRISTOPHER K. RIESBECK

SUPPORTED BY
NATIONAL INSTITUTE OF MENTAL HEALTH
AND
ADVANCED RESEARCH PROJECTS AGENCY
ARPA ORDER NO. 457

FEBRUARY 1972

COMPUTER SCIENCE DEPARTMENT
School of Humanities and Sciences
STANFORD UNIVERSITY



PRIMITIVE CONCEPTS UNDERLYING VERBS OF THOUGHT

BY


Roger C. **Schank**, Neil Goldman
Charles J. Rieger and Christopher K. Riesbeck

ABSTRACT: In order to create conceptual structures that will uniquely and unambiguously represent the meaning of an utterance, it is necessary to establish 'primitive' underlying actions and states into which verbs can be mapped. This paper presents analyses of the most common mental verbs in terms of such primitive actions and states. In order to represent the way people speak about their mental processes, it was necessary to add to the usual ideas of memory structure the notion of Immediate Memory. It is then argued that there are only three primitive mental ACTs.

This research is supported by Grant PHS MH 06645-09 from the National Institute of Mental Health and (in part) by the Advanced Research Projects Agency of the Office of the Secretary of Defense (SD-183).

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

Reproduced in the USA. Available from the Clearinghouse for Federal Scientific and Technical Information. Springfield, Virginia 22151.
Price: full size copy, \$3.00; microfiche copy, \$.95.



PREFACE

This paper grew out of an ongoing seminar at the Artificial Intelligence Project on machine understanding of natural language. It became clear to participants at that seminar that a set of primitive actions must be established **in** order to create adequate conceptual structures. We have debated about the most common verbs in English, attempting to fit them into existing primitives. Whenever we could not do this we created a new action primitive. If we could not break down this primitive we let it stand.

The participants of this seminar, which is still going on at this writing date, are: David Brill, John Caddy, Neil Goldman, Linda Hemphill, Chuck Rieger, Chris Riesbeck and Roger **Schank**. All of the above people have contributed to the work presented here. The authors of this paper were responsible for the actual writing of the paper but many of the ideas were arrived at jointly during the seminar.



TABLE OF CONTENTS

	<u>Page</u>
1. Introduction	1
2. Conceptual Dependency	3
2.1 Conceptualization	3
2.2 Paraphrase Recognition	10
2.3 Summary	12
2.4 Physical Primitives	13
3. Conceptualize	14
4 ... MTRANS	20
5. Belief	24
5.1 Believe	24
5.2 Want and Like	32
5.3 Knowledge	38
5.4 MTRANS vs PTRANS	41
5.5 Behavior Beliefs	43
6. Immediate Memory	48
7. Verbs of Mental Combination	56
8. Important Extras	70
8.1 Understand	70
8.2 MTRANS vs PTRANS again	76
8.3 Judging	79
8.4 Volition, Intention, and Their Relations to Actions	83
9. Conclusion	94



1. Introduction

For the past four years there has been an effort undertaken at Stanford to enable computers to understand natural language sufficiently well so as to be able to perform in a dialogue situation. We have attempted to analyze natural language into meaning structures that are unambiguous representations of the meaning of an input utterance. We have required of those representations that they be unique. That is, the meaning representations of any two utterances which can be said to convey the same meaning should be identical.

Thus, we have concerned ourselves with the creation of conceptual structures, and the predictions and inferences that are possible given a formally defined conceptual structure.

The initial form of a conceptual dependency structure was intended to be a language-free unambiguous representation of the meaning of an utterance. In fact, the conceptual structures that were initially used (Schank (1969a) and Schank (1969b)), bear a great deal more similarity to the surface properties of English than we now believe should exist in such structures. Subsequently, we began looking for common concepts that could be used for representing the meaning of English sentences, that would facilitate paraphrase by the conceptual structures without losing information. The concept '**trans**' was introduced (Schank, Tesler and Weber (1970)) as a generic concept into which words such as 'give' and 'take' could be mapped, such that by specifying attributes of the cases of 'trans' no information **would** be lost. (For example, 'trans' where the actor and recipient are the same is realized as the verb 'take', whereas, the actor and donor part of the recipient case are

the same, the verb is 'give'). Such generic concepts simplified the conceptual networks, making them more useful. Furthermore, it became apparent that the linguists' problem of the representation of such concepts as 'buy' and 'sell' became solvable. Semanticists such as Katz (1967) have argued that while these concepts seem close enough it would be arbitrary to choose one as the basic form of the other, so the correct thing to do must be to write formal rules translating structures using 'buy' into structures using 'sell' when this is deemed necessary. Instead of doing this, we made the suggestion (Schank (1970)) that using 'trans' one could **map** 'buy' into 'trans money causes trans object' and 'sell' into 'trans object causes trans money'. Such a representation eliminates the 'which is more primitive than the other' problem and instead relates the two events that actually occurred.

The naturalness of the concept 'trans' led us to consider whether there might, in fact, be more of these generic concepts around. Thus we began a search for primitive concepts that can be used as the basis of conceptual structures. This report discusses the results that we have arrived at. In order to appreciate them however, it will be necessary to set out the rudiments of the conceptual dependency framework first. We shall present in the next section the basics of conceptual dependency as well as some important changes that have occurred since the last A.I. Memo (Schank (1971a)).

2. Conceptual Dependency


2.1 Conceptualizations

In previous papers, we have stated that we were using what is basically an actor-action-object framework **that includes** cases of the actions. We wish to make clear that that framework is precisely what we are using, and furthermore we are using it quite literally. That is, any action that we posit must be an actual action that can be performed on **some** object by an actor. Nothing else qualifies as an action and thus as a basic ACT primitive. The only actors that are allowed in this schema are animate. That is, an action is something that is done by an actor to an object. (The **exception** to this rule regards natural forces which shall not be discussed here.)

Actors, actions and objects in our conceptual schema must correspond to real world actors, actions and objects. To illustrate what is meant by this consider the verb 'hurt' as used in 'John hurt Mary'. To treat this sentence conceptually as (actor: John; action: hurt; object: Mary) violates the rule that conceptual actions must correspond to real world actions. 'Hurt' here is a resultant state of Mary. It does not refer to any action that actually occurred, but rather to the result of the action that actually occurred. Furthermore, the action that can be said to have caused this 'hurt' is unknown. In order to represent, in our conceptual structure, an accurate picture of what is going on here the following conceptual relationships must be accounted for: John did something; Mary was hurt; the action caused the resultant state. In conceptual dependency representation, actor-action complexes are indicated

by \Leftrightarrow , denoting a mutual dependency between actor and action; **object-**state complexes **are** indicated by \Leftrightarrow denoting a predication of an attribute of an object or by \Leftrightarrow denoting a change of state in the object: Causal relationships are indicated by \Uparrow between the causer action and the caused action, **denoting** a temporal dependency. Causal arrows (\Uparrow) may only exist between two-way dependencies (\Leftrightarrow , \Leftrightarrow or \Leftrightarrow). That is to say, only events or states can cause events or states.

Thus our representation for this sentence is:



The dummy 'do' represents an unknown action. 'Hurt' is ambiguous between mental hurt (**hurt_{MENT}**) and physical hurt (**hurt_{PHYS}**).

Conceptual dependency representation then, seeks to depict the actual conceptual relationships that are implicit within a natural language utterance.


Actions, in conceptual dependency, are things that are done to objects. Actions sometimes have directions (either through space or between humans), and always have means (instruments). These things are called the conceptual cases of an action. Unlike syntactic cases, (as posited by Fillmore (1968) for example) conceptual cases are part of a given action and therefore are always present whenever that action is present. Thus, if an action takes an object, whether or not that object was mentioned it is considered to be present conceptually. If the particular instance of that object was not stated and is not inferable then an empty

object slot is retained.

The conceptual cases are: OBJECTIVE; RECIPIENT; DIRECTIVE; and INSTRUMENTAL,. The first action we shall introduce is 'ptrans' which is used in any physical transfer sentence. The sentence:


John gave Mary a book.

is conceptually analyzed using 'ptrans' and objective and recipient case as follows:



The symbol \leftarrow^o denotes 'object of the ACT' and the symbol \leftarrow^R denotes 'recipient of the object', with the recipient of the object in the 'to' part, and 'donor of the object' in the 'from' part.

Actually, this analysis is not quite correct for this sentence since the sentence is conceptually ambiguous. The conceptual diagram above is correct for one sense of the sentence but it is possible that the transition was not done physically by John. Rather, John could have said 'you can have the book' and Mary could have taken it herself. Since we don't know what specifically John may have done we represent this sense as:



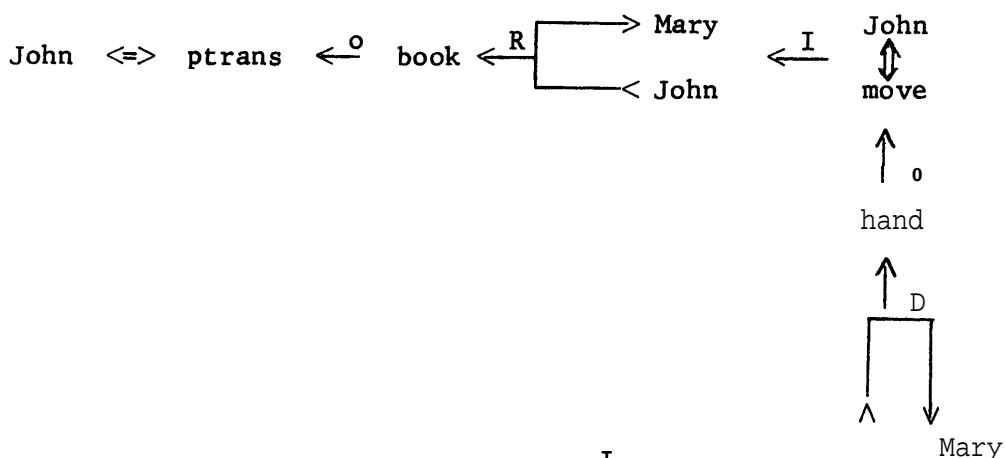
Either of these two structures may have been the intended one, but we assume unless given information to the contrary that the first is correct.

Suppose the sentence had been:

John gave Mary a book by handing it to her.

Here, the sentence is disambiguated by the 'by clause'. All actions

require an instrument that is itself another actor-action-object complex (called a conceptualization). When the action in the main conceptualization is known, it is possible to delimit the set of possible instrumental actions. For 'ptrans' the ACT that is most often the instrument is 'move'. 'Move' represents the physical motion of a **bodypart** (which may be holding an object) by an actor, together with the direction that that action takes. The conceptual analysis of (3) then is:



The instrumental case is indicated by \xleftarrow{I} and the conceptualization that is the instrument is dependent upon (written perpendicular to) the main conceptualization. The directive case (indicated by $\xleftarrow{D} \left\{ \begin{array}{l} \rightarrow \\ \leftarrow \end{array} \right.$) shows the physical direction of the action. Thus 'the book was moved towards Mary'. (It is necessary to indicate here that the hand is holding the book also, but we shall not enter into that here.)

Since every ACT has an instrumental conceptualization that can be said to be part of that ACT, we can see that it should therefore be impossible to ever actually finish conceptually diagramming a given sentence. That is, every ACT has an instrument which has an ACT which has an instrument and so on. In this sentence we might have conceptually something like: "John transed the book to Mary by moving the book towards Mary by moving his hand

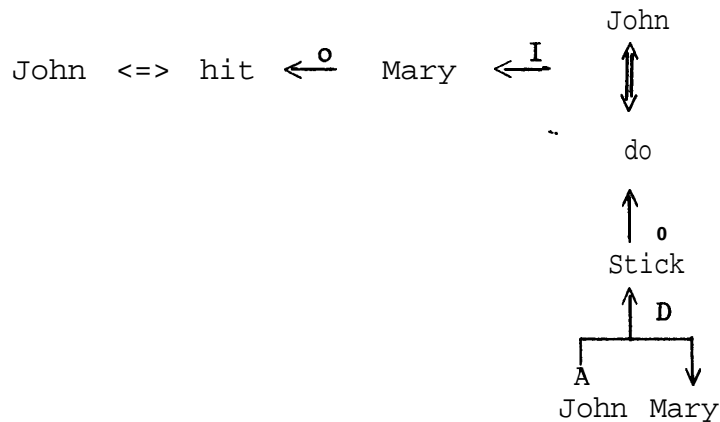
which contained the book towards Mary by grasping the book by moving his hand towards the book by moving his hand moving muscles by thinking about moving his muscles" and so on. Since an analysis of this kind is not particularly useful and is quite bothersome to write, we do not do so. Rather, whenever we represent a conceptualization we only diagram the main conceptualization and such instrumental conceptualizations as might be necessary to illustrate whatever point we are making. It is, however, quite possible that we might need many of these instrumental conceptualizations in a program that was intended to simulate certain body motions (such as Winograd's (1971) block moving program). Thus, the ACT in a conceptualization is really the name of a set of actions that it subsumes (and are considered to be a part of it). These instrumental conceptualizations are not causally related since they are not actually separable from each other. In actuality, they express one event and thus are considered to be part of one conceptualization. The rule is then, that one conceptualization (which may have many conceptualizations as a part of it) is considered to be representative of one event.

In ordinary English usage, the syntactic instrument of a given sentence corresponds conceptually to either one of two potential places in a conceptualization. Either it represents the object of an instrumental conceptualization (usually the first instrumental conceptualization) or it is the object of a conceptualization that causes the conceptualization most directly related to the verb of which it is an instrument syntactically. Conceptually an instrument can never be only a physical object. Thus as an illustration of the first instance we have:

John hit Mary with a stick.

We represent the conceptual action underlying 'hit' by 'hit' meaning 'forceful

physical contact'. Thus we have conceptually:



The 'do' in the instrumental conceptualization indicates that the action by which the physical contact was done is unknown. This corresponds to the fact that this sentence is actually ambiguous. The two most **common** interpretations being that 'he **swung** the stick' or that 'he threw the stick'. Representing such a sentence in this manner allows for the discovery of this ambiguity. (In an actual computer analysis schema the blank 'do's' can be realized as predictions about missing information which must be discovered either by inquiry or memory search.)


Predictions about what ACT's fit into this instrumental slot are made from the ACT in the main conceptualization. 'Hit' requires either 'move' or 'propel' as actions for its first instrument. 'Swing' and 'throw' are mapped conceptually into 'move' and 'propel' respectively (with additional information as to manner).


The other type of conceptual realization for a syntactic instrument can be illustrated by:

John grew the plants with fertilizer.

Traditionally, linguists would consider 'fertilizer' to be an instrument of the verb 'grow'. Conceptually however, 'grow' is simply a state change


and is not an action that can be performed by someone on something else. Rather, a person can do something that effects this state change. Thus we have as the basis of the underlying conceptualization:



The 'do' in this conceptualization represents the extremely important fact **that** something was done by John. Thus the plants were not 'grewed', they **grew**. (represented by  for state change). What John did was not 'causing', rather what he did caused something else to happen.

Since the 'do' represents an unknown action, it might be of interest to find out what that action might have been. But since that information was unstated, finding it is the job of any processor that uses the **results** of a conceptual analysis.

The syntactic instrument of 'grow' is treated conceptually then as the object of the causing action. Thus we have:



We can, in fact, make an educated guess as to what John could have done with fertilizer that would have caused the growing. Probably he moved it to the ground **where** the seeds were. Since this is an inference we shall only mention it here without going into how to figure out such a thing.

2.2 Paraphrase Recognition


Before going on into the substance of this paper, it might be interesting to consider how such a deep conceptual analysis of natural language utterances can help us in parsing and understanding those utterances:

Consider:


John prevented Bill from leaving the room.

The verb 'prevent' is conceptually a statement about the relationship of two events, namely that one event causes the inability of the occurrence of a second event. Unless we treat 'prevent' in this manner, important paraphrase recognition ability will be lost, and in addition even the ability to intelligently parse sentence derivative from this will be hindered.

Conceptually then, 'prevent' is not something that anyone can do, rather it expresses the following relationship between two events.



That is, person₁ doing something caused **person₂** to not be able to (c) do something else. Thus we have:




If we had an intelligence understanding system, we might want to know what John 'did' and this representation allows us to realize that we could ask that. We might also want to know where Bill intended to go, but that is less likely. Now consider:

John prevented Bill's leaving the room by hitting him.

Along with the information that 'prevent' represents the conceptual structure

shown above is a clue as to how to go about finding what might fill in the first 'do'. This clue is that if the ACT that replaces the 'do' is present it is most probably in the syntactic instrument of 'prevent', that is, in a by-clause.

Thus, that clue is used to give us:



It is important to notice that it is quite possible to realize the above structure as the following sentences as well.

Bill couldn't--leave the room because John hit him.

When John hit Bill it caused Bill to be unable to leave the room.

When John hit Bill, it meant that Bill had to stay in the room.

The above sentences do not use 'prevent' in word but they do use the concept underlying 'prevent'. It is extremely important that any theory of understanding analyze these sentences or any of the myriad other paraphrases into only one conceptual structure in a natural way. This requires establishing the relationships between actual events rather than between the words that may have been used to describe those events. In order to do this, it is necessary to break words down into the primitive actions and events that they describe. That is what we seek to do in this paper with respect to the mental verbs.

2.3 Summary

In summary then, conceptual dependency is a representation for expressing the conceptual relationships that underlie linguistic expressions. The basic structure of this conceptual level is the conceptualization. A conceptualization consists of either an actor-action-object construction or an object-state construction. If an action is present then the cases of that action are always present. One case of an action is instrumental which is itself a conceptualization.

Conceptualizations may be related to other conceptualizations causally. Just as it is impossible to have an action without an actor, so it is impossible to have the cause of a conceptualization be anything other than another conceptualization. (This means that 'John moved the table' must be conceptually, 'John did something which caused the table to be in a different **position**'. This doing is not 'move' but rather something that was unstated. The doing can be inferred and is most probably 'apply a force to'.)

Other requirements on conceptual relations are not stated here because they would only complicate matters. **Schank** (1972) is a good source for those. There are, however, more relationships that we shall use here for which we have not introduced notation. If one object stands in relationship to another as possessor (x possesses y), we express this conceptually as:

If such a relationship is predicated by a sentence we indicate it as: $y \iff \text{POSS}(x)$ or for Location: $y \iff \text{LOC}(x)$ (y is located at x). Here LOC indicates the type of state predication (\iff).

Thus, if we had 'John is in my house', conceptually it would be:

John \iff LOC (house)

A POSS
I

2.4 Physical Primitives


Throughout this paper we shall be concerned with conceptually representing verbs that have to do with mental things. Since we shall be **compelled** to use examples in our discussion that pertain to the physical world, we shall introduce here the ACT primitives of the physical world that we use. We make no attempt to justify the **physical ACT** primitives here, the current plan being to do so in a forthcoming paper.

We use the following **ACTs** to describe the physical world, from which most physical verbs can be derived:

<u>ACT</u>	<u>Required Cases</u>	<u>Actors</u>	<u>Objects</u>	<u>Meaning</u>
ptrans	O,R,I	human	physobj	change receivership
move	O,D,I	animate	bodypart	move bodypart
ingest	O,D,I	animate	food	ingest
hit	O,D,I	animate	animate/physobj	forceful contact
propel	O,D,I	animate/n.f	physobj	apply force to
go	O,D,I	animate	animate	move oneself
Look-at	O,I	animate	physobj	
smell	O,I	animate	physobj	
Listen-to	O,I	animate	sound	
speak	O,I	animate	sound	
grasp	O,D,I	animate	physobj	
physcont	O,I	animate	physobj	contact

These eleven primitive actions are used to describe the physical world.

Now we shall begin to discuss the representation of mental world phenomena.



3. Conceptualize

The first mental activity we shall consider is that denoted by the English verb 'think', in the sense of 'thinking-about' concepts or ideas. It is a process fundamental to activities described by many English verbs, and thus an excellent candidate for inclusion in any set of primitive conceptual ACTS.

We postulate the existence of a primitive ACT, CONC, which refers to the act of conceptualization. We shall write:

$$\text{ACTOR} \Leftrightarrow \text{CONC} \xleftarrow{\circ} X$$

Where X may be any conceptualization. We further restrict CONC to human ACTORS (although one might argue that some animals engage in 'conceptualizing' as well). Since conceptual dependency theory hypothesizes that people think on a conceptual, rather than a linguistic level, and since people frequently receive and transmit information about 'thinking', the ACT CONC is needed to provide a conceptual representation of this activity and of linguistic information pertaining to it. The requirement that the **object** of CONC be a conceptualization is merely a consequence of the above mentioned hypothesis.

The ACT CONC is that which in English is referred to as 'to think-about' in a very broad sense. By CONC we mean:

- i) to focus attention on, as well as
- ii) to perform mental processing on, where mental processing may include finding associations, and may, through another mental ACT called MBUILD (described in Section 7), result in implications, inferences, etc.

If CONC involves these other processes then it is reasonable to inquire why it should be considered a primitive ACT. The answer involves understanding the notion of a conceptual primitive as distinguished from an empirical

primitive. . Just as those **ACTs** which we hypothesize as physical conceptual primitives (such as MOVE-ing a bodypart) may be described in greater detail (e.g., by neural activity and muscle contraction) so may the ACT CONC be described as a complex of operations on an information store in any given instance. These operations are of the type used in numerous computer implemented memory models (e.g., Quillian(1966) and Becker (1969)) and involve manipulation of various links and nodes. The reasons for not breaking CONC into these more primitive terms are just as strong as (and analogous to) the reasons for our choice of physical conceptual primitives:

- (A) Even if we were able to specify a relatively small group of truly primitive mental processes which cover the activity we now **call** 'conceptualizing', there would be no way to know which of these processes was being referred to when a verb such as 'think about' was encountered.
- (B) These more primitive processes appear to have no more definitive meaning than does CONC except in relation to a reasonably sophisticated memory model. And we do not wish to assert that a human being refers to any such complex memory in using natural language. Thus we choose CONC as a conceptual primitive.

[Note that (A) and (B) are not independent observations. Given that humans are not aware of models incorporating some set of psychologically primitive mentalactivities, they cannot be directly differentiating these activities in the language which they use. This, in turn, means that no parser should be expected to understand this language in terms of these primitives.]


One other point should be mentioned in regard to the use of CONC. It is true that whenever a person speaks he has **CONC-ed** the conceptualization

which represents the meaning of his utterance. We do not, however, wish to represent this CONC-ing act as a part of the meaning of that utterance. CONC will be used only when the utterance itself refers to certain mental activities, which may have been performed by the speaker or another person. (A similar verbal action, 'entertain' is posited by Price (1969)).


Following is a representative sample of English 'mental activity' verbs and senses in which they can be described conceptually by CONC:

THINK - ABOUT

"John is thinking about eating an apple."



"Mary is thinking about John."



(Here Cl refers to some particular conceptualization, and INVOLV indicates that the dependent appears in Cl.)


We are maintaining the requirement of the conceptual syntax that the object of CONC be a conceptualization, not a concept. Although the syntactic object of the verb 'think-about' may be a noun, we claim it is impossible to conceptualize the isolated meaning of that noun. One may only conceptualize a conceptualization in which that noun fills some role. If we do not know what that conceptualization is, we must represent it with a dummy of some sort.

Thus if Mary tells someone "I was thinking about John yesterday", the

presence of the element C1 in the representation of this utterance might lead him to ask her "What was it you thought about John?". In other words, the representation makes it very clear that Mary hasn't really said what she was thinking about.


DREAM

"Bill dreamed he was a doctor."



CONSIDER (one sense)

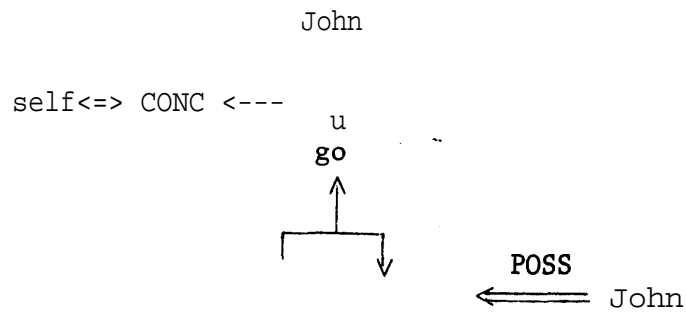
"John considered going home."



Here no distinction has been made between 'consider' and the first sense of 'think - about'. The difference seems to be that when we hear 'consider' we expect the act to result in the ACTOR's making a decision. But another way-of viewing this is to say that English speakers choose 'consider' in those cases in which the object of the conceptualizing is a future action or state over which the 'conceptualizer' has **some** control. Thus, while it is perfectly understandable, most English speakers would not say: "I considered having wasted two hours yesterday", but rather "I thought about having wasted two hours yesterday".

WONDER

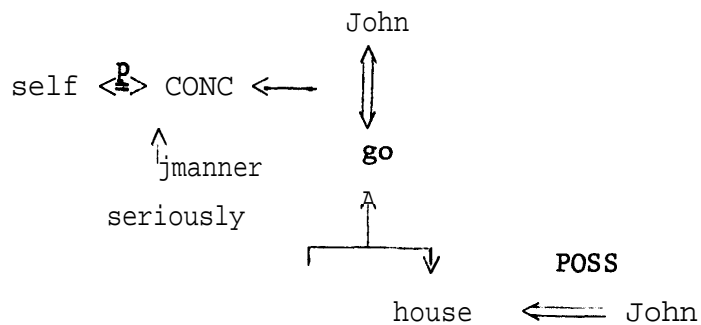
"I wonder if John is going home."



The point here is that the verb 'wonder' indicates CONC with an object conceptualization having the question (?) aspect indicating that the relationship between 'John' and 'go' may not have occurred.

PONDER

"I pondered John's going home."



This example gives rise to the question of which English verbs or phrases represent MANNER modifications of CONC. For example, the idiom 'to give passing thought to' seems to fall into this category.

While the English language certainly modifies 'conceptualize' verbs with manner adverbials, it is not pleasing to do so conceptually. Such a representation requires the use of modifiers dependent on CONC and it is not clear what sort of modifiers of actions, if any, are conceptually possible here. It is rather dubious whether something like 'seriously' can be considered a conceptual primitive of any type at all.

Another solution to this problem is to modify CONC with respect to

time. Actions do have duration and this will have to be representable in conceptual notation. It is plausible that CONC-ing manner adverbials can be handled by duration modifications. 'To ponder' or 'concentrate on' then means to conceptualize something for a period considerably longer than the norm, while to 'give passing thought to' requires the opposite sort of modification.

4. MTRANS

Once we have the action 'conceptualize', we must consider that it is necessary to do certain actions in order to conceptualize and furthermore that people talk about such actions. That is, given that there is a representation for something being in memory, the problem of how to handle the simple and basic actions of bringing something from and putting something into that memory comes next. The act MTRANS described below is meant to handle this basic flow of information to and from the conscious mind.

It, plus various mental building acts, should serve to represent all the ways in which we bring thoughts into our heads.

MTRANS:

MTRANS represents a change in the mental control of a conceptualization (or conceptualizations) and underlies verbs like recall, commit to memory, perceive, sense, and communicate. It has several features different from the physical TRANS. For one, the object that is **TRANSed** does not leave control of the donor, but is copied into the control of the recipient. Further, the donor and recipient are not two different people but two different mental processors (or locations: the distinction in the mind is as fuzzy as the distinction between program and data in the computer), which are frequently within the same person. Five such processors will be used here:


1. Conscious Processor (CP) - this operates on concepts that one has become aware of, performing deductions, making choices, forming associations, and other such actions.
2. Long Term Memory (LTM) - this is primarily the store of beliefs one has about the world. It is a processor too, where such actions as forgetting and subconscious association occur, but the level of activity is both low and hard to characterize, so it shall be

treated as a passive element here.


3. Immediate Memory (IM) - this is like the LTM and is meant to represent the short term event memory humans use to keep track of propositions relevant to the current situational context.
4. Sense-Organs (Eye, Ear, Nose, Tongue, and Skin) - these are all pre-processors, converting raw sense data into conceptualizations describing that data.
5. Body - this covers whatever processors handle internal sensations, such as pain, unease, excitement, etc.

With these items, we can handle many mental verbs, such as


I remembered **Bill** was a communist:



I saw Mary sleeping:



I feel pain:



This use of MTRANS covers mental actions where the concept brought into awareness has been internally arrived at, rather than externally generated.


Verbs that refer to externally generated conceptualizations include:

COMMUNICATE:




This is pure communication, mind to mind, i.e., telepathy. With the instrumental case to modify the means of communication we can represent more mundane, indirect verbs like:

I told him Mary was asleep:




Forgetting is simply the inability to bring something from LTM:



Verbs such as 'learn' and 'teach' also involve MTRANS to LTM from CP.

Thus:

I was taught that Bill was a communist.



That is, 'teach' is really like communicate. The actual difference lies in the fact that the communicated information is said to be new in the case of 'teach'. Thus, we also have the information that this information was not in the LTM of self before. In order to represent this we shall have to

discuss 'believe' (Section 5).

Nothing has **been** suggested in the above as to what the mental acts are that create new concepts. Hopefully, however, it has been shown that with one simple action which we've called **MTRANS** it is possible to handle a noticeable number of mental verbs, **including verbs** of communication, reasonably straightforwardly. Various necessary instrumental acts involving the physical component of perception - e.g., looking at something, making vocal noises - have not been discussed on the assumption that they, like many body actions, are primitive for our purposes.

5. Belief

5.1 Believe

In modelling a human language user, it is necessary to model his belief structure, since both understanding and generation of natural language require recourse to such information. Many lexical items refer, directly or indirectly, to a **conceptualization's** presence in or absence from some individual's belief structure. Conceptual dependency theory must, then, be able to represent the relationship between a belief and a person's 'world model'.


For the purpose of explaining the verb 'believe' and other related verbs, it will be necessary to use our previous partition of the human mind. First, we claim that the objects which people manipulate in the process of thinking are conceptualizations. Conceptualizations which are available to be acted upon by CONC must be located in the individual's CP. All other conceptualizations are stored in LTM or **IM**.

(The notion of a mental location is not to be taken literally. One must consider "in CP" to be a property which some conceptualizations may have at some times. This property could be expressed as "having activation tag" if we wished to avoid the location analog, but any phrase we choose will have erroneous associations if taken literally, so we shall stick with the notion of mental locations. We are not claiming the existence of separate sections of the brain with graphs being shunted in and out of them. What we are claiming is a conceptual reality for the existence of these three distinct properties of conceptualization which are present in the mind. We will denote mental location by MLOC.)

Finally, we claim that all conceptualizations in X's LTM are held to be true by X. Thus if X does not believe that Y is a doctor, he will not have in his LTM the graph

Y \Leftrightarrow DOCTOR

although he may have such conceptualizations as:



or..

Y \Leftrightarrow DOCTOR

The English verb, 'believe' has two primary senses. The first of these is paraphrasable as "hold to be true". We will represent this sense by:

x \Leftrightarrow MLOC (LTM(Y))

This represents "Y holds X to be true", where X may be any conceptualization and Y a conceptual nominal with the property "human" X may represent a simple observation about the world ("The ball is red."), a philosophical viewpoint ("Fascists kill babies.") or a rule of behavior ("If person1 hits **person2**, person2 should run.").

The question now arises as to just what is meant when someone says "Y believes that X ". If Y is speaker (i.e., "I believe that...") the above representation seems legitimate. But what if the speaker is not Y? Isn't it somewhat too limited a representation?

It is clear that users of English have different criteria for evaluating the truth of the statement "Y believes X". Hardly anyone requires that Y

have X stored in one particular form in his belief structure in order that this statement be true. For example, if Y has stored the beliefs:

- a) John is 6' tall, and
- b) Mary is 5' tall

then almost anyone would accept that Y believes

c) John is taller than Mary,
even though a simple deduction on Y's part is necessary to arrive at this belief.

On the other hand, if X is a belief which would require a complex chain of reasoning on the part of Y, involving perhaps deduction, association, and analogy, we might answer the question "Does Y believe X?" with "Y has probably never thought about it", or just plain "No", even though we think that Y would believe X if he were asked and if he thought about it long enough.

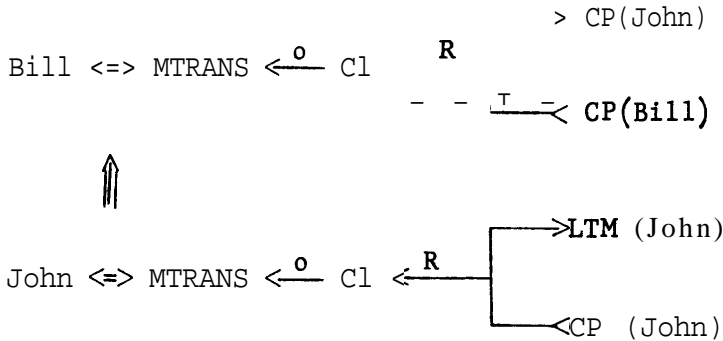
The fact that humans use the term 'believe' in slightly different ways does not seem to create havoc when the term is used, and it seems unlikely that we create an internal model of what 'believe' means to each person with whom we communicate. We therefore see no need to create an operational definition of 'believe' before using this as a primitive concept. Any reasonable computational model will contain such a definition **implicitly in** its program -- a definition utilized whenever the model is asked "Do you believe...".

At any rate, if one keeps in mind the distinction between a physical reality for LTM, which we do not claim, and the conceptual reality of a property of a conceptualization, which we do claim, then no confusion should result from the use of the above representation for both utterances "I

believe..." and "John believes....".

The second sense of 'believe' involves an ACT which results in a person's believing something. We represent this sense using the ACT **MTRANS**. For example:


"John believes Bill."



where C1 represents any unspecified conceptualization. From this it is 'deducible that:

$$\text{C1} \Leftrightarrow \text{MLOC}(\text{LTM}(\text{John}))$$

Note that we often have ambiguity in the past tense. "John believed *what Bill said", may also be:



One final point should be made before we move on to some examples. Just as people normally CONC what they say, so too do they normally believe

what they say. Nonetheless, the fact that the speaker believes X is not what he means when he says X, but merely an inference from his having said it. And when we are interested in representing the meaning of an utterance we must keep such inferences distinct.

A brief comment on notation is appropriate at this point. While we would write:


$$Y \Leftrightarrow *BELIEVE* \xleftarrow{o} X$$

to represent that Y believes X, we feel that representing *BELIEVE* as a primitive ACT is misleading. What we refer to by the first sense of 'believe' involves no actual action, mental or physical, but is better understood as revealing a static membership property of the specified belief. Thus our representation is a notation more consistent with the semantics of conceptual dependency than the above.

Following is a representative list of English lexical items and senses that can be described conceptually in the above manner,


AGREE

"John agrees that Bill is guilty."



The meaning of "Y agrees that X" is the same as that of "Y believes X". 'Agree', however, contains an additional message to the parser (hearer) that X is a conceptualization which has occurred previously as a belief.


"John agreed with what Bill said."



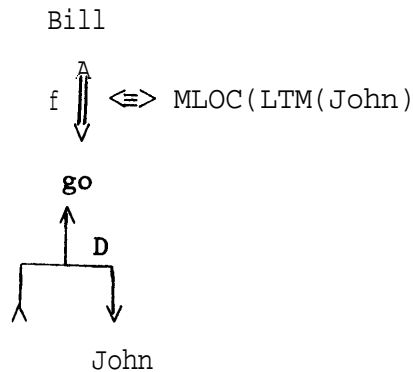
In this sense of 'agree' we are given the original source of the conceptualization and thus the parser would be required to find the original conceptualization rather than its source.

EXPECT

a) "John expects Bill to **become** a doctor."



b) "John expects Bill."




Both these senses of 'expect' convey the same meaning as 'believe', but this meaning is realized as 'expect' only when the belief is a future action or state.

Above we asserted that all conceptualizations in X's LTM are held to be identically true by X. This is a very strong assertion and is not what people really appear to do. More importantly, people convey through language information which indicates that they do not believe they do this.

We can increase the explanatory power of the above notation if we use the notion of a credibility weight tagged onto a stored conceptualization. We will indicate this in diagrams by a new property (CRED) of a conceptualization. For the purposes of this paper we shall assume only credibilities of low, medium, and high.


BE (SURE, CONFIDENT) THAT

"John is sure that Bill is guilty."




SUSPECT

"John suspects that Bill is guilty."



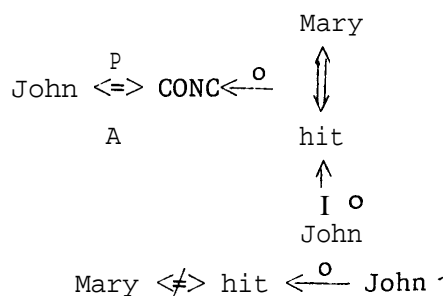
DOUBT

"John doubts that Bill is guilty."



IMAGINE

"John imagined that Mary hit him."



FEAR

"John fears bears."

Bears \Leftrightarrow do



\Leftrightarrow MLOC(LTM(John))

John \Leftrightarrow hurt

'Fear' always indicates a belief by the 'fearer' of the existence of a causal relationship between some action and injury to the 'fearer'. In cases where the action is specific, (often expressed by 'fear that') it often indicates, in addition, a belief that the action will take place.

There are other senses of 'fear', among which is one in which the English object appears as an instrument conceptually rather than as an ACTOR -- e.g., "John fears guns."

5.2 WANT AND LIKE

The three verbs 'want', 'like', and 'please' appear at first to be only distantly related to the mental activities being discussed in this paper. However, the distinctions between the three turn out to be inextricably linked to both the concepts of CONC-ing and believing (MLOC = LTM). Of course, when one considers that 'want' is at least very close to being the opposite of 'fear' this fact should be somewhat less surprising.

Before embarking on a detailed analysis of these words let us point out two pitfalls which may have caused considerable confusion in the discussion before:


- (i) We must distinguish the case in which the source of an utterance involving one of these verbs is also the want-er (liker-er, one who is pleased) from the case in which the source is a second party.
- (ii) The tense of want (like, please) as well as the tense of the object clause may reveal components of the meaning.

Consider the sentence "I like to eat ice cream." The fact that I am saying this enables a listener to infer:

- (a) that I am CONC-ing it, and
- (b) that I believe it.

but these inferences should not necessarily be construed as part of the meaning of my utterance. What I definitely am communicating is, first of all, that eating ice cream has put me in a pleased state at some past time, and, in addition, that in general (timeless) eating ice cream has this

effect on me. Thus we have:



Note that one does not use the present tense of 'like' with future tense objects - "I like to go to Europe next year" - although one does use 'want' in such cases (more of this later).


When 'like' takes its past tense form - "I liked eating ice cream" - we again can use the above representation, but with the timeless marker (Δ) deleted.

Consider now the more revealing case of a second party source - "John likes to eat ice cream". Certainly the information conveyed above with 'self' replaced by 'John', is again being communicated. But now we cannot even infer that John, who may be asleep sanewhere when this statement is made, is CONC-ing anything about ice cream. What about John's beliefs? Normally we can infer that he believes eating ice cream leaves him pleased. But this is only true because the ultimate source of information about an internal state of a person X is almost always X himself. Thus we reason that if someone tells us that John likes something, John must have been the ultimate source of that information and thus must believe it himself. But consider the statement: "Chuck doesn't believe it, but he likes having his wife squeeze his big toe while he sleeps." This statement claims that someone likes something without believing that it pleases him. If we can accept this as a normal use of the verb 'like', and don't claim that it is a second sense of the verb, then the belief

component must be an inference and not part of the meaning of 'like'. [In some dialects the above sentence would sound somewhat more natural with 'enjoy' or 'please' than with 'like'. If the sentence seems to be contradictory in **some** dialect it may be that the belief aspect has actually become part of the meaning of 'like']. "

Finally, consider the use of '**like**' in the future tense: "John (will, would) like seeing that movie." Here no claim whatever can be made about John's beliefs - he may have never heard of the movie. If we consider 'like' to have a belief component in the present and past tenses then we must consider this to be a new sense of the verb. Otherwise the **same** conceptual structure can serve all three tenses.

Let us now perform a similar analysis of the verb 'want'. When someone says he wants to do something, he is saying he believes doing it will cause him to be pleased. Thus the 'do-cause-pleased' structure is present in 'want' just as it was in 'like'. But no claim is being made that the pleasurable activity was ever done before, nor that it will always bring pleasure. This points out the first basic difference between the two verbs. The object clause of 'want' is always future tense (with respect to the time of wanting). If the time is not specified explicitly we can generally insert 'now' meaning 'in the immediate future'. Thus for "I want to eat ice cream" we have at least the following:




In this case, with the source of the statement being the want-er, we again have the fact that the source both CONC-ed and believed at the time of the utterance, but it is **not** clear whether this is part of the meaning

of 'want' or a valid inference based on the source of the utterance.

In "John wants to go to the movie next week" we have the case of a second party source of an utterance. As with 'like', it is certainly not the case that the speaker is communicating anything about what John is currently CONC-ing. But neither is he saying that going to the movie next week will please John. (The speaker may be of the opinion that John will detest the movie if he sees it). What he is communicating is a belief of John's, in particular:

..




John \Leftrightarrow pleased


In fact, this representation gives the intuitively plausible paraphrase from "John wants..." to "John believes he would like..."

In summary, we have the following basic structures to represent the verbs 'like' and 'want' and 'please':

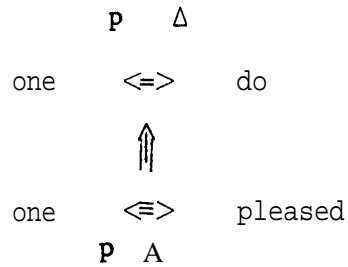
LIKE



WANT



PLEASE



which gives the often discussed paraphrase from "X likes Y" to "Y pleases X." Those who distinguish these phrases probably have a belief component in their usage of 'like', in which case 'like' and 'want' look very similar -- except for their tense markers.

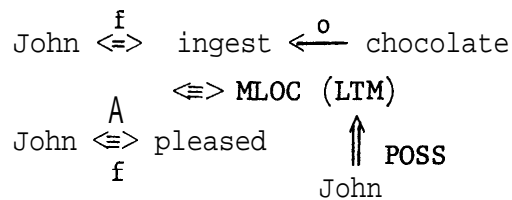
Another general aspect of the verb 'want' which is not restricted to the discussion--of mental verbs alone, concerns its different syntactic senses. Sentences such as:

"John wants some chocolate", and

"John wants his mother"

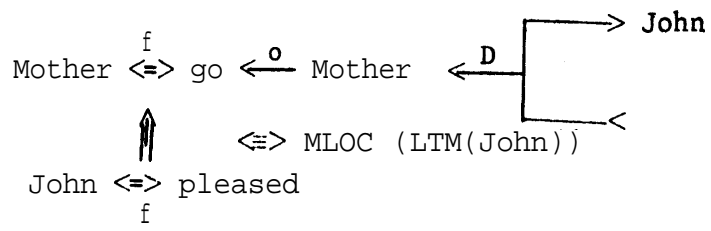
require a parser to make inferences to decide what actually is the conceptualization which John believes will please him. Once this conceptualization is found it is clear how it is to be entered into the 'want' structure to represent the meaning of such a sentence. In this case we have:

"John wants some chocolate."



Here an inference has been made that John wants to eat chocolate, but for

"John wants his mother"




5.3 Knowledge

The question of what is referred to as 'knowledge' and what is 'belief' is a sticky one which has been the subject of much discussion. We have seen no solution to this problem which appears entirely adequate for conceptual dependency theory, and we do not claim that the brief treatment given below is without its own problems.

A simple way to dispense with this problem is to claim that the distinction between knowledge and belief is just a matter of degree of -credibility; in other words, 'knowing' is just 'being very sure'. Somewhat surprisingly, this does handle satisfactorily a great deal of the usage of 'know', and should certainly be listed as one sense of 'know' in our dictionary. The discussion which follows should be seen more as an attempt to elicit the conceptual distinctions between the verbs than as an argument for a particular conceptual representation.

KNOW THAT

a) "John knows that Bill is guilty."



Here we claim that "John knows that X", where X is a proposition, is equivalent to 'John believes that X, which is true'. The element *TRUE* in


the above representation refers not to any absolute sense of truth, nor even to any logical sense, but rather to the fact that the proposition is believed not just by John, but by some other person or persons (usually the source of the statement 'John knows that X'). Thus when the utterance is encountered, two pieces of information are being received:

- 1) John believes X, and
- 2) the source of the utterance believes that X is a widely held or empirically verifiable proposition.


KNOW (IF, WHETHER)

a) "John knows whether Bill is guilty."

i) $Bill \Leftrightarrow guilty$



Bill



\wedge

ii) $TVAL = TVAL2$

where $TVAL1$, $TVAL2$ are variables which may take on the values TRUE or FALSE.


Thus to 'know whether X' is to hold a belief which assigns the correct

• (from the viewpoint of the source of the utterance) truth value to X.

The ambiguous statement "John doesn't know whether Bill is guilty" has the following two meanings:

b) John holds no belief as to Bill's guilt.

Bill \Leftrightarrow guilty



This C-diagram is obtained by negating the main link of (a.1). We are assuming here that TVAL may range over its entire domain (TRUE,FALSE).


c) John holds an incorrect belief as to Bill's guilt.

i) A (TVAL1≠TVAL2)


Our representation of 'know whether' thus provides an explanation of the source of ambiguity which arises when **the** verb is negated -- namely, that the meaning is represented by two structures, either of which may be negated separately.

5.4 MTRANS vs PTRANS


Now that we have discussed 'believe' it is possible to make clearer the analogy between MTRANS and PTRANS. MTRANS behaves like PTRANS in that we can talk about coming to believe something (coming into possession of something) versus believing something (having something). That is,



is like




While "I believe that."




is like "I have it."



Hence, to represent "I believe he will go." we write



This is not to be confused with the sense of feeling as associated with sensations, such as



'Hoping', which is a verb in English, is handled here as a state, as a special case of believing. To 'hope' is to 'believe' something good has the possibility of happening.

"I hope to go."

self <=> CO

↑
possible

A

↑↓
MLOC(LTM(self))

self<=> CO

↑
<=> PLEASSED

It should be noted here that 'possible' is not related to the way the world is, but the way we see the world. That is, not only are future events possible, but so are events that have already occurred but about which we have not heard news. Logically a past event has either happened or not, but we can say (and have understood) sentences like "It's possible he went yesterday," and "I hope I passed." In English "Event p is possible" means "I do not know and cannot deduce the fact that p did not occur."

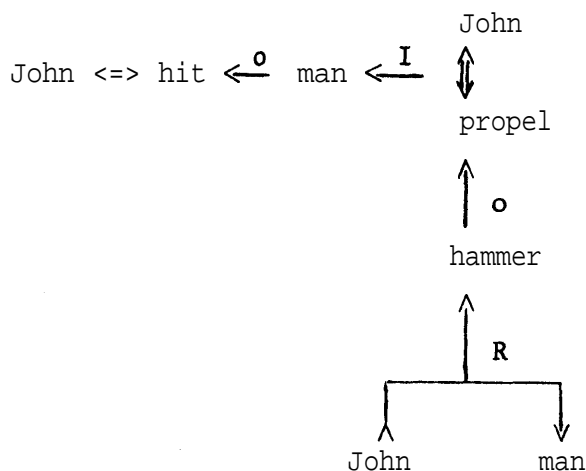
5.5 Behavior Beliefs

Since we are discussing belief it might be worthwhile to digress a moment to show that in sane cases belief patterns (i.e. the 'should' and 'ought to' types of beliefs) are often expressed by means of words other than believe. Adverbs in English are a prime example. Consider the following sentence;

(1) "**John** hit the man with a hammer with vengeance."

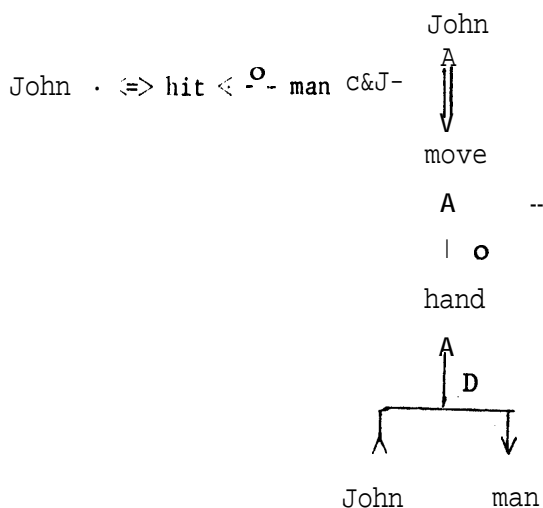
Conceptually this sentence has two principal senses. In one, the man was hit by a hammer that was thrown at him. In the other, the hammer was swung at him. The conceptual structure underlying the first sense of (1) represents the fact that 'John propelled the hammer in the direction of the man' and that this conceptualization - event was responsible for the contact of the hammer and the man.

This is written in a conceptual dependency structure as:



The second sense of sentence (1) means that John was holding **a hammer** and moved it in the direction of the man such as to effect a contact of the hammer and the man. This sense is written in conceptual dependency as:

John $\langle \Rightarrow \rangle$ grasp $\langle \overset{O}{\leftarrow}$ hammer
 A



Now the interesting question is, how does a concept such as 'vengeance' fit here? First, it is obvious that even though 'vengeance' is a noun in English, semantically it would seem to be a paraphrase of 'vengefully' and thus, semantically at least, modifies the verb. But what is its function conceptually?


If we were to simply modify the ACTS involved (move and propel) we could not simply attach vengefully as a modifier of these sections. Conceptually that explains nothing. The only possible modifiers of primitive conceptual actions are those that actually refer to aspects of those actions. -Consider 'move'. We have said that the primitive action 'move' is used whenever a **bodypart** is 'moved'. Clearly, the only kinds of modifications of such motion are those of path travelled and speed. That is, the only variant types of 'move' are things like move quickly, move steadily, move with acceleration, move in a swinging fashion, move directly, move with a chopping motion, and so on.

The question is then, can vengeance be a description of the speed or path of a moving object? Since it cannot (the only sense in which it could, belongs to the realm of inference), we have to find some other place for it.

It is important to realize that 'vengeance' is simply another form of 'revenge'. In order to deal with a meaning analysis of the concepts of a sentence containing 'vengeance', it is necessary to deal with the meaning of 'revenge'. 'Revenge' is not a simple word by any means. The reason for this is that 'revenge' and 'vengeance' are expressing a belief pattern. Thus in order to analyze (1) correctly we shall have to correlate it with the belief pattern that is expressed within it.

We define a belief pattern as a prescription for action that expresses a value on the part of the speaker. That is, the kind of beliefs of which we are speaking are of the form 'if X happens then one should do Y', or 'an X is one who is likely to do Y', or 'X is bad' and so on.

With respect to sentence (1) 'vengeance' can be said to be reflective of the following belief-conceptual structure:




The 'cf' on the causal link refers to the situation that something may cause something else. This structure represents that if person one does something which results in person two being hurt in some way this could cause person two to do something that is intended to result in person one being hurt in some way. This belief is labeled in English as 'revenge'. It is what speakers of English understand by the word 'revenge' even if they themselves do not

believe that such a response is justified given the initial condition. The structure given is simply that elicited by the word 'revenge'. The word 'vengeance' calls this structure as well. Moreover, when the word 'vengeance' is presented, the conceptualization underlying the sentence that 'vengeance' modifies can be placed in the 'one₂ \Leftrightarrow do' part of the above belief. That is, it was this conceptualization that **was** done in response to some previous hurt in the view of the speaker. Thus, the speaker is saying that the hitting of the man appeared to be in response to an act done by him that hurt John. This statement by the speaker has nothing to do with the actual truth or falsity of such an assertion.

Thus we are saying that an accurate dictionary entry for the above words would read as follows:

vengeance, revenge, vengefully, revengefully:

Call following belief:



The conceptualization that is modified by the word under discussion in the sentence is to be placed in slot (3) in the above belief.

Belief patterns such as this one also exist in the memory of the speaker and many English adverbs refer to them. Such belief patterns serve as reasons for given actions (usually (3)).



1

.

...

6. Immediate Memory

We **hypothesize** an Immediate Memory to account for certain abilities of humans that could not be accounted for without it. Roughly, Immediate Memory (IM) can be considered to contain those items from Long Term Memory that are currently being used or are likely to be used by the Conscious Processor.

If you are in a room, talking to some people about their lives, and then decide to play bridge, IM will contain information about these people, your relationship to them, and about bridge; both your experience with bridge and the rules and vocabulary of the game.

It is IM that takes care of language processing. That is, one can only speak about those words and concepts that are in IM. Anaphoric reference is only possible for items that are in IM. If an item is referenced that is not in IM confusion will result and either nothing will be retrieved as referent or a time lag will occur while an item is retrieved from LTM. Thus, the IM reranks the senses of word with respect to the context. If the word 'ruff' is used in a bridge context (or the word 'bridge' for that matter), the correct concept associated with that word will be chosen without reference to any possible dogs named 'Ruff' (or any other kind of 'bridge'). That is, the IM functions as a sort of context keeper that reorganizes what is at the ready surface of memory. If 'card' meaning 'postcard' were used in the bridge context, there is no reason to expect that the sentence would not cause sane confusion; particularly if the context of the sentence in which 'card' appeared did not serve to point out the intended sense, e.g. 'I didn't like getting thatcard'.

Thus IM serves to store readily available information. Such information


is erased from IM when the context changes. Some contexts, such as facts about the world that are needed every day or individual personal contexts, never change. Such things are always in **IM**. Thus IM has no specific time period after which there is decay, but rather the decay is instantaneous and entirely dependent upon context.

Thus, we would need to talk about an IM in any model of human **memory** simply because such a thing appears to exist. Also any efficient language analysis algorithm must make use of the contextual information contained in **IM** so as not to make useless exhaustive searches of spaces that are severely limited by context.

It is the *case* that people talk about IM and we must be able to represent this. IM is responsible for new information that is entered into a human memory before it is deposited in LTM. Not all new information is stored (remembered) in LTM, but such information can be retrieved for a while after it is outlined. New information then, is stored in IM and only sometimes transferred to LTM. If the context changes and the transfer has not been made, such information is forgotten.


Thus the store of beliefs which people seem to use, those that are in the Immediate Memory, is unlike the LTM store in that it varies from situation to situation. It contains both new beliefs that are accepted for the moment and old beliefs that seem relevant to the current environment. Some of these beliefs may be inconsistent with other beliefs in the LTM but unless dissonant beliefs are brought together their inconsistency will not normally be noticed. Part of LTM then consists of clusters of beliefs appropriate to various situations, that are transferred as a block from the LTM to the IM. (See Yinger (1965) for a discussion of situational clusters.)

One very common example of beliefs that are accepted for the moment are those assumptions made in a novel that is being read or a movie that is being watched. Although it could be said that such temporary beliefs are actually permanent ones with the context specified -- e.g., "in this book this twenty-six year old girl dies" -- still the fact that people cry at movies and boo the villain indicates that to some extent the false world takes on a practical reality. IM is a place where such suspension of disbelief can safely occur. This momentary believing, expressible in English as "accepting for the moment" is represented as:




In addition, IM takes care of things that we believe for the moment but which have little long term value.

"I'm sure I believed him at the time but I didn't bother to remember it."



A



In addition, things that affect us, affect our emotional state, and must be either in the CP, in which case we are aware of them, or in the IM, in which case we may only be aware that something is on our minds.


"Something's been bothering me."

concept \Leftrightarrow MLOC(IM(I))




Things that keep impinging on our awareness, whether emotionally charged like an insult, or emotionally neutral like a simple melody, reside in the IM, the neighbor to the CP.

"I keep remembering what happened."



Taking one's mind off something is not a simple process of **MTRANSing** the concept out of IM. MTRANS is not an actual moving action in the sense that PTRANS (physical trans) is; that is, just because something has been **MTRANSed** from X doesn't mean that it no longer exists in X. MTRANS means that it (or perhaps a copy) also exists in another location. This is reflected in the fact that one cannot be ordered to forget something. Forgetting is not something one does, but something one allows to happen. To have forgotten something is to be unable to retrieve it.


"I forget whether Bill is a philatelist or not."



Hence to get something off one's mind is to preclude the transfer of the troubling item from IM to CP.


"She took his mind off the problem."

she <=> do

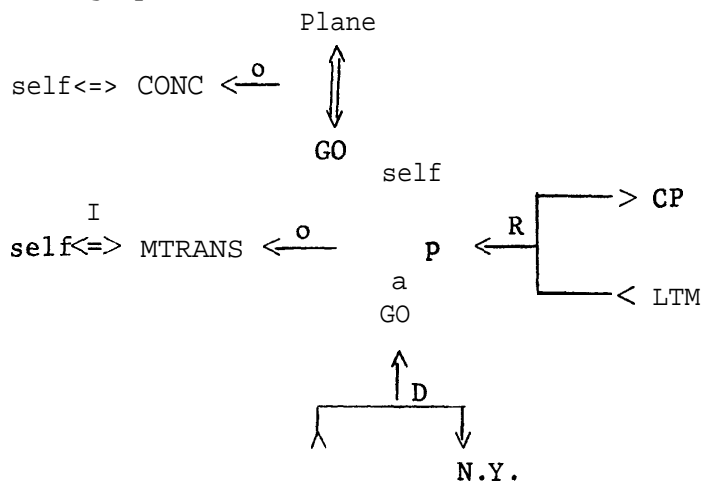


It is possible for one to think of things that are not in IM, of course. The most obvious example of **this is** the word 'remind'. "Remind" is represented by a causal link between thinking about (CONC) something (usually just perceived) and remembering (MTRANS) something else. The general structure for representing "remind" then is:

- "X reminds one of Y!"



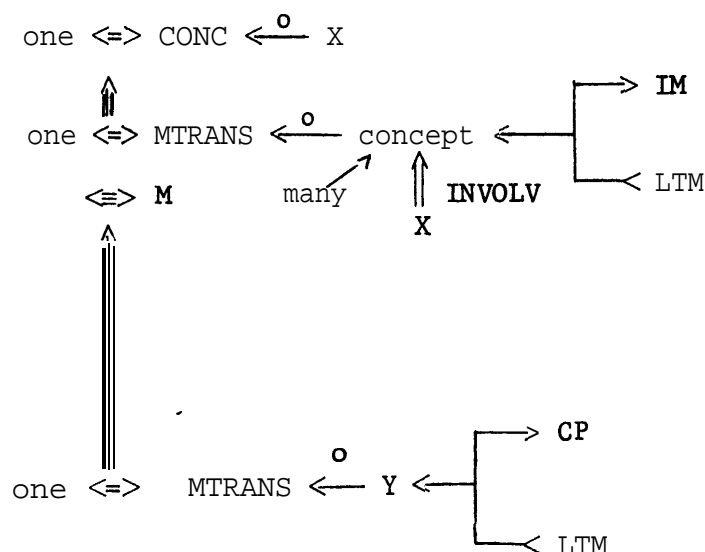
For example, "That plane taking off reminds me of the time I went to New York," is graphed:



A point that needs to be justified here is that if CONC and MTRANS are to be used consistently, their objects, X and Y, must be conceptualizations. When X and Y are events, **as above**, this doesn't seem unreasonable, but what

about sentences like "Your dog reminds me of my husband?" The problem is related to the whole question of what is the nature of objects of mental acts. The claim is that nothing less than a conceptualization, i.e., a complete thought, exists by itself. **Both** in the mind and in the world, everything is imbedded in some context. In the physical world, any object is involved with relationships of location, contact, force, and so on. In the mental world objects join with those mental relationships which we notate as CD links like, \Leftrightarrow , $\xleftarrow{\circ}$, $\xleftarrow{\left[\begin{array}{l} \rightarrow \\ \leftarrow \end{array} \right]}$, \Leftarrow , and so on. The primacy of--context is similar to Quillian's (1966), but rather than looking at the mind as a net of words, it is seen as a general net of special nets of conceptual primitives, i.e., as a net of conceptualizations. When we think about an object, we are thinking about, or have on our mind (in IM), a cluster of conceptualizations about that object. "Your dog" is a set of facts and memories, like "your dog has short hair," and "Your dog has a nasty bite." In becoming aware through perception of some of these facts, one causes to be brought into IM those other things one knows about "your dog" - those things that define him for the speaker. Some of these facts, which may not even be in the awareness of the speaker, through some associative Linking item, cause facts about "my husband" to be brought into awareness.

A fuller representation of "X reminds one of Y", then, might be:

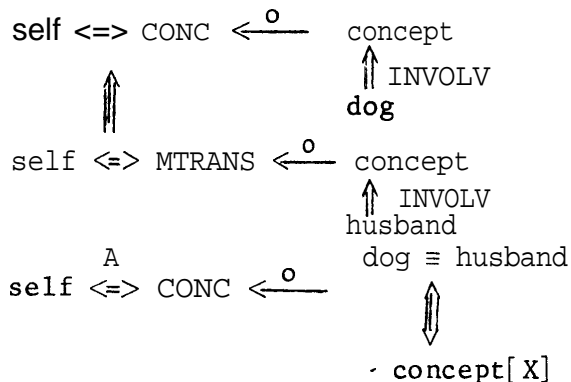


where X and Y are conceptualizations, perhaps unspecified, as in
 "Your dog" which is



The middle conceptualization though is an inference that can be drawn when someone says "Your dog reminds me of my husband, though I can't think why." The essential factors in the phenomenon of "reminding" that people refer to in language, however, are the thoughts we were aware of that preceded the new thoughts, and these factors are what we include in our definition of "remind".


"Your dog reminds me of my husband " has a bit more to it than we have yet represented. This is the sense of "remind" that is discussed by Postal (1970), the sense meaning "Your dog and my husband seem similar in some way." That is, the concept I had of your dog has brought up some concept of my husband which leads me to note a similarity. In CD representation, we can say that $X \equiv Y$ (X is equivalent to Y) with respect to $\langle \Rightarrow \rangle$ some conceptualization.



There are, of course, restrictions on what conceptualization pairs can lead to a conclusion of equivalence. If 'concept-about-husband' equalled 'concept-about-dog' in form, but for the substitution of "husband" for "dog" everywhere, and if "concept" were not a trivial statement, such as "dog \Leftrightarrow ingest \leftarrow food," then we could state a conclusion of similarity. Given just the equality in forms we could still conclude similarity, but we would not bother to state it.

For example:

"Your dog is like my husband in that they both have short hair."

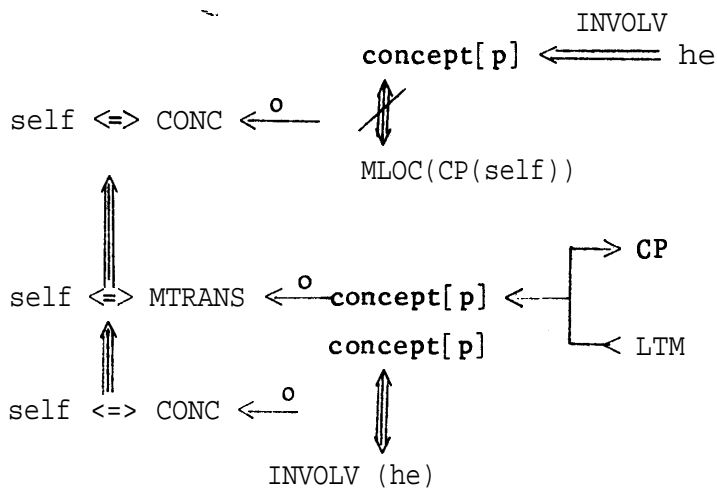


where X is a place-holder for forming true conceptualizations about "dog" and "husband".

Two concepts can be sufficiently related to lead to a conclusion of similarity even without this strong equality, however. For instance, "your dog is like my husband in that they both run fast" involves two concepts with different instrumentals for the going, namely four legs and two legs, respectively. This is a minor difference between the concepts, however, and predictable from what we know about people and dogs. Such obvious modifications are automatically made by people. It will be necessary at some point to have a conversational program capable of such obvious changes as well. The use of a surface verb similarity between two predications that does not carry over onto the conceptual level is as bad a violation

of the reasonability criteria for \equiv as using two unrelated statements. The incongruity does form the basis for a set of jokes, however, such as "Her teeth are like stars; they come out at night."

'Recognize' is 'remind' plus. We recognize X when thinking about X (usually upon perceiving it) reminds us of some situation involving X which we met before. For example, I recognize John when seeing him causes me to remember a previous time when I met him, or saw his picture, or heard a description of him. Recognize then has the property of using memory structures to attach relationships to some input. Recognizing is the attaching of relationships to some object within the **conceptualization** (e.g., answering "Who left?").



The notation "Concept[p]" indicates that the conceptualizations involved have a past, "p", tense modification. That is, thinking about someone I can't place (i.e., I have no previous knowledge - **concept[p]** - involving him in CP), reminds me of some facts that I realize do involve him. This realization (the CONC) is the point at which recognition occurs, although as soon as the proper memories (concept[p]) are brought in, the realization follows immediately.



7. Verbs of Mental Combination

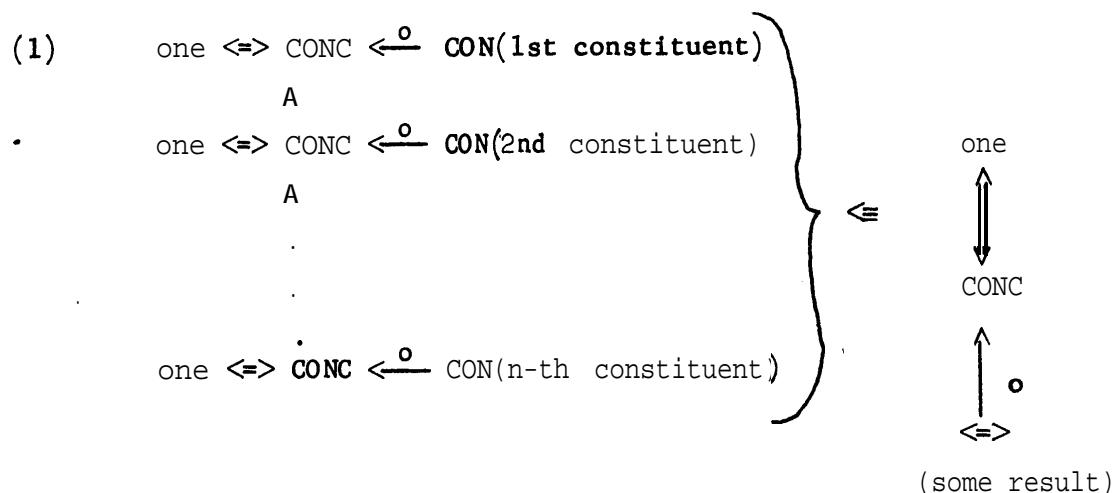
In the same way that we build a complex physical structure from many building-block constituents, related only insofar as they will each contribute in **some** prescribed way to the final product, we also build new mental structures as assemblages of conceptualizations. We piece together conceptualizations as though sewing together a quilt, and the end result is both more than the "sum" of its parts and undeniably different from each contributing part. This process of conceptual combination and integration is an important mental process: were we only to experience conceptualization after unrelated conceptualization, as though viewing each frame of a motion picture individually, we could never extract any order, logic or continuity from what we perceive or conceptualize.

Mental combination has many manifestations. For instance, we can "put two and two together", coming up with some new conceptualization; we can extract differences or recognize the generalities among several conceptualizations; conceptualizations can interact in the form of poorly understood imagery; we can engage in feats of logical deduction; we can "weigh the evidence", and arrive at a conclusion, by however a circuitous route. The list of our abilities is seemingly interminable when expressed in these terms. All of these processes are commonly called "thinking".

The outstanding questions before us are these: (a) do such processes represent primitive mental actions?, and (b) if so, is there one primitive mental act which can account for the lot, or are these processes symptomatic of a very involved set of primitive actions? (Both questions are of course relative to our fairly restricted scope of inquiry).

The first question is the easier to answer in light of the development of the subject represented by this paper. We have so far posited the existence of two other closely related primitive acts: CONCEPTUALIZE, AND MTRANS, as well as "exist in LTM". To answer the question of whether there are mental combination primitives is to determine whether or not these existing primitives can capture in any way the notion of conceptual combination. "Exist in LTM" and MTRANS clearly have no expressive power in this respect: the former is a statement of a static condition, and applies only to one conceptualization, however complex it may happen to be; the latter expresses only the movement of mental objects: it specifies the flow of the stream, not the confluence of two streams to produce one.


It is a slightly more involved task to demonstrate CONCEPTUALIZE to be independent from acts of mental combination. In particular, there is at least one deceptively attractive method of expressing mental combination in terms of CONCEPTUALIZE and appropriate graph structure:



However, there are at least two problems with this approach. First, it violates the definition and intent of CONCEPTUALIZE, which is the

representation of that portion of cognition which is the focal point of all mental activity. It assumes a singularity of attention, a uniquely total allocation of cognitive resources. As such, there can only be one conceptualization active at a time, even if it is a very complex one. (For instance (1) in its entirety could conceivably be CONCEPTUALIZED once it has been pieced together and "loaded" into the conscious processor. But this is not the argument at hand. The question is: how did the pieces come to be part of the same conceptualization in the first place?; what is the additive process of conceptual bonding?).

We could view the process in a more serial way:

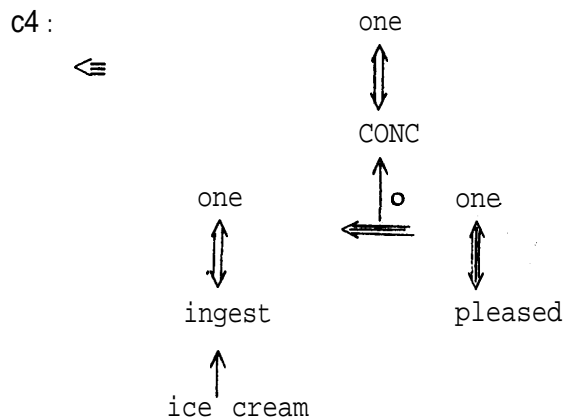





but we are still faced with the same dilemma. Either each **MTRANS** "overwrites" the contents of the conscious processor established by the preceding **MTRANS** (the motion picture problem), or there is an "**MTRANS-and-add-on**" process, which is the conceptual combining agent. Even if we ignore this bonding problem for the time being, there is something still misdirected about permitting the focal point of the conscious processor to service more than one conceptualization at a time. Such a thing as cognitive focus does exist, even if it is only introspectively detectable. When we speak of combining conceptualizations, we are in fact indicating that we have chosen one of them as central to the process, and are conceptualizing just that one. Then, from somewhere on the periphery of the conscious processor,

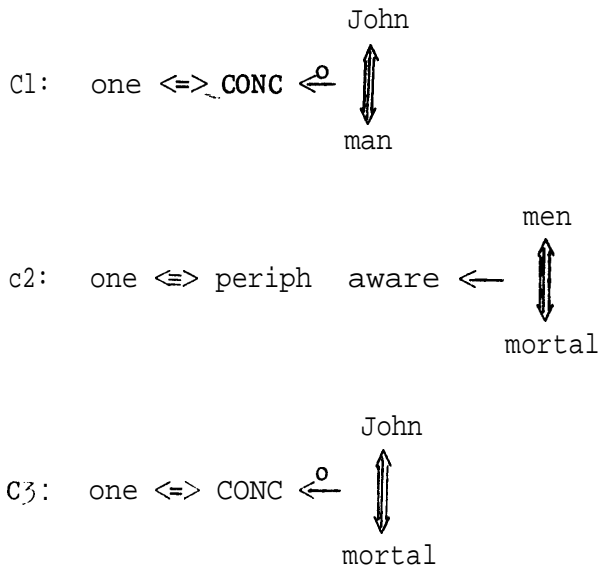
other conceptualizations "lying in wait", "jump out" and combine into the conscious processor. This is a far more spontaneous process than either approach (1) or (2) would imply. Furthermore, it implies the existence of some peripheral awareness: that limbo of conceptualizations which is everything left over after excluding the conceptual processor contents **and all** "inactive" contents of LTM.

(This peripheral awareness is just that partition of cognitive storage which is termed immediate memory (IM) elsewhere in this paper. It is, **roughly**, all information connected by associations to that information currently in the conscious processor. As such, it is totally transient. It may alternatively be viewed as that fleeting "active" subset of LTM peripherally related both to the contents of the conscious processor and to "recent" contents of the conscious processor.)

The second problem stemming from using CONC in this way is subtler than the first, and uncovers the real flaw: regardless of how CONCEPTUALIZE (or **MTRANS** to CP) is strung together in CD graphs in an attempt to account for the combination process, the (often) highly constrained and structured relation between the constituent conceptualizations is only incidentally implied, not explicitly mentioned. We might just as well write things like



whenever we need to represent mental combinations. Absurd as such a graph seems, we do not want to rule out such extreme cases of "associative spontaneity". It is perhaps the case that in P's head C1, C2 and C3 do indeed elicit C4, and further that all are necessary to elicit C4. Whatever the associative mechanism at work here might be, it unquestionably exists. The point is that this type of combination is truly best served by the causal structure of (3), since the constituent conceptualizations were never really combined, and were certainly not "related" in any other sense. On the other hand, form (3) is not adequate for expressing such highly structured processes as



where C3 is the conclusion reached using C1, C2.


Thus it is clear that, while certain cases of mental combination expressible in terms of CONCEPTUALIZE, MTRANS and "exist in LTM" do exist, these three primitives will not suffice for more structured forms of mental combination. We are left with the job of ascertaining how many primitives are required. There are two ways to approach this problem: (a) to relate mental processes to the processes of another domain by way of analogy, and (b) to examine and explain as many mental verbs as possible using some

primitive or a set of primitives. We shall pursue both avenues.

There seems to be a highly parallel system of actions in the physical domain: those acts which surround the primitive notion of putting **some-**thing together, of assembling the sum from the parts, of causing the existence of an object. Such an object, on the one hand, exists before its creation in the substance of each individual part, yet on the other hand, does not exist, lacking the relations governing the assembly of its constituents. In the mental domain, we encounter precisely the same paradox: our brain is a gold mine (!) of uncombined constituents (conceptualizations). Yet most of them will never surface in the right groupings, or in the right context, to produce the potentially infinite range of novel ideas (conceptualizations). The new ideas exist in their potential, yet do not exist in reality.

The notion of creating the sum from its parts is a very **general** one, being applicable to everything about us: the motorcycle is built from parts, the tree is constructed from definable units, love is the sum of various behaviors, which in turn are products of physiological events, hydrogen is composed of atoms, atoms of sub-atomic particles, and so on. The point is that to "build" something, to generate something new out of old parts, is a very specific concept abstractly, yet very non-specific in its domain of applicability. We would argue that "mental build" has precisely the same attributes of specificity yet generality within its mental domain. However, whereas physical building verbs are usually masking variable underlying actions, for mental building the underlying actions ("micro-actions") are not variable for our purposes, but are characterized by a specific mental ACT. We, therefore, have reason to believe that one primitive can explain many processes of mental combination.

Now we turn to practicalities: what English verbs correspond directly to or are expressible in terms of this primitive, which we will christen MBUILD and write as



(Note here the introduction of a new CD notation. We use the many-to-one 'functional' arrow to denote the combination and transformation of several units into one resultant unit.)

The words of the following list are intimately related to MBUILD, and will hopefully clarify its nature:

- | | |
|--------------|---|
| "think over" | ("I'll think it over.") |
| "consider" | ("I'll consider all the facts.") |
| "deduce" | ("I deduced that the butler did it.") |
| "reason" | ("I reasoned that if a,b,c then d...") |
| "conclude" | ("I concluded that we ought to lock up all commies .") |
| "compare" | ("Now class, today we'll compare Brazil with Idaho.") |
| "prove" | ("I proved to myself that it would work.") |
| "resolve" | ("I resolved the problem of how to get home.") |
| "solve" | ("I solved the problem of how to get home.") |
| "relate" | ("I'm trying to relate what you said to this.") |

Before getting into specific examples, notice that in some of these verbs, MBUILD plays the role of the action which is antecedent to some more "final" act of accepting the result as knowledge or as a belief. Examples of this type are "conclude", "resolve", "prove to oneself", "solve" and so on.

In these cases, an end result is actually produced and its CONCEPTUALIZATION is therefore implicit. In others of these, MBUILD is the only ACT underlying the verb, and there is no result conceptualization yet produced (such as "think over", "consider", "reason out", "relate", etc .) This distinction between the process and the result of the process (and what becomes of the result afterward) is crucial to the unravelling of mental verbs. MBUILD refers only to the process of combination, or attempted combination, and includes no information about the success or failure of the operation. Success can be denoted by the presence of a result in the CD graph notation, failure by its absence.

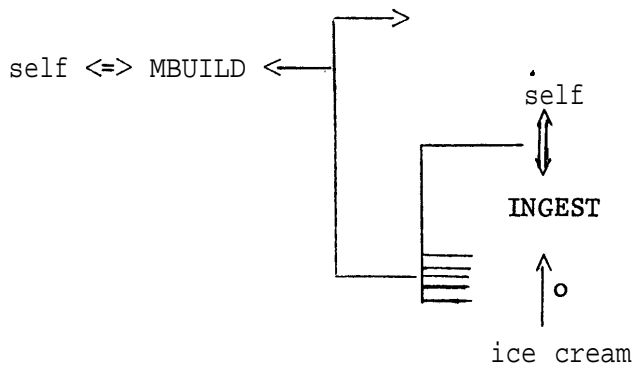
Another point to be made is that the contributing conceptualizations may or may not be made explicit in their role as "arguments" to MBUILD. Quite often, in fact, they are not made explicit, or only some of them are. This is more of a notational variation of **MBUILD's** use, although missing constituents could serve the useful purpose of providing motivation and direction to an "understanding" program using this primitive. As with other forms of missing information in CD graphs, these unknowns could initiate memory searching and reasoning processes in hopes of filling them **.in**, or if those procedures failed, could generate a relevant question (MAN: I concluded that.. .; MACHINE: What makes you think that?), Similar remarks apply to unspecified result positions (MAN: I've been thinking about these things ; MACHINE: And what have you concluded?).

It should **be clear** that **MBUILD's** are often related, and expressed, sequentially. If we ask the mad scientist how he discovered X, he will probably tell us: "First I realized that a,b, and c were the case. I reasoned, therefore, that d followed from these. Then I discovered e which,

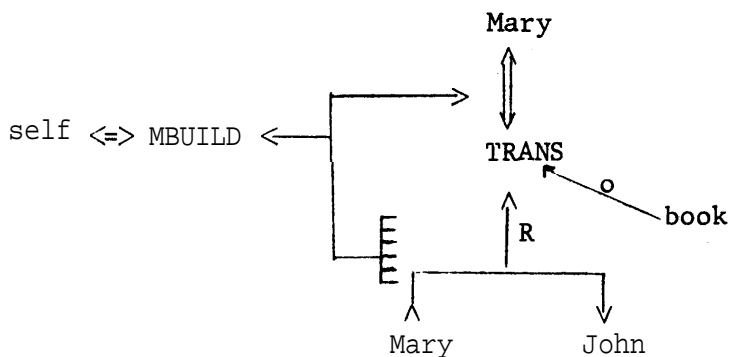
in light of d, can only mean f, . . ." In other words, the result of one MBUILD quite often becomes a constituent (**antecedant**) of another MBUILD.

EXAMPLES:

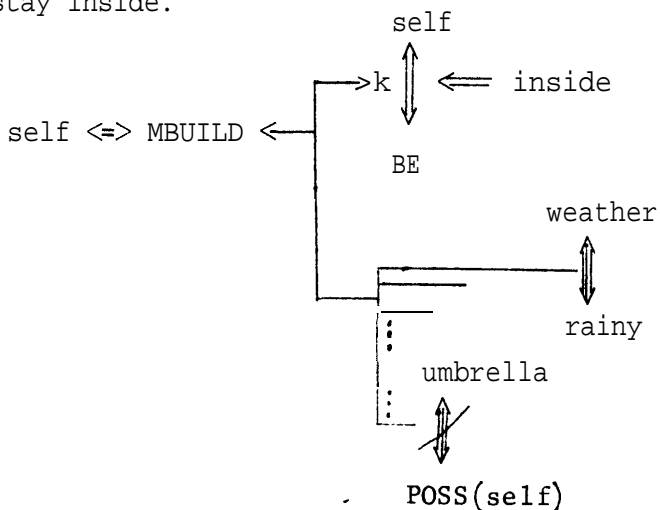
I'm considering the ramifications of eating that ice cream:




I concluded-that Mary gave John the book.




Since it was rainy and I had no umbrella, I figured that I ought to stay inside.




I realize that these facts a and b are unrelated.



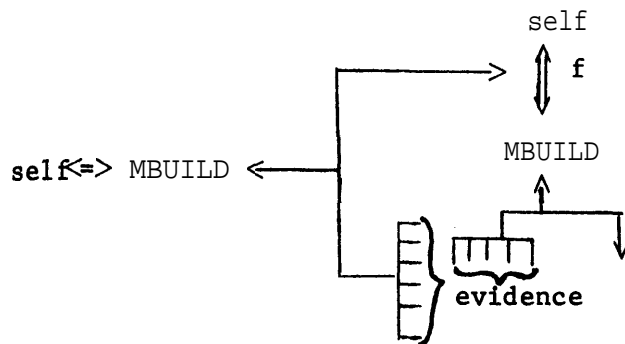
I won't even consider these facts a and b.



I convinced myself that it was unnecessary to go.




I have weighed the evidence and decided to reconsider.




Have you thought about the problem (P) yet?

you $\stackrel{p?}{\Leftrightarrow}$ MBUILD




What did you conclude?

you $\stackrel{P}{\Leftrightarrow}$ MBUILD



Why did you conclude c?

you \Leftrightarrow MBUILD

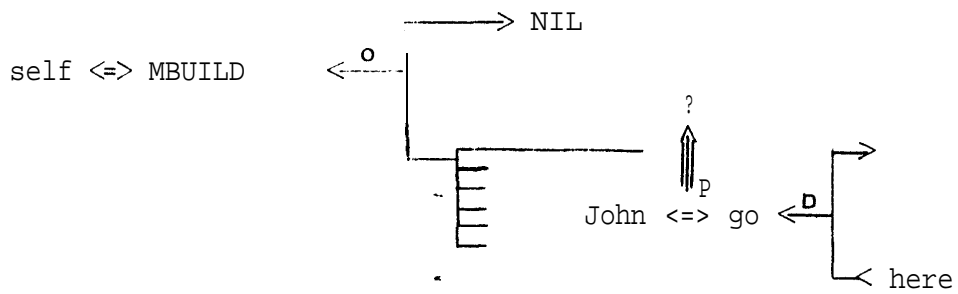


There is one further clarification to be made regarding the relationship of the arguments of MBUILD to the MBUILDing process. There are two cases which we have lumped together in the examples: a) the MBUILDing occurs in "free-form" (is non-directed), and b) the MBUILDing is "directed" by one of its arguments. The first case is characterized by the paradigm: "Here are some things to think about. What can you conclude from them?" In this case, there is no particular problem in mind to direct or constrain the MBUILD to one domain. The second case is that of finding the solution to a.. particular problem, the answer to a particular question. In this case not only is the MBUILD process "directed" by the problem, but the kinds of other arguments MBUILD will use are implicitly "related" to the problem. Perhaps these two cases actually represent quite different mental and logical processes. Yet MBUILD seems to be central to both, and their differences involve **"micro-processes"** which we do not need for the purposes of CD.


How do we notate directed MBUILD? During the course of answering a question, we are aware of the question itself. To this extent, the question itself is not only directing the MBUILD, but is also one of the arguments of the process. There is a direct analogy here to theorem proving by the resolution method: to answer a question (prove a theorem), the negation of the question is resolved against other facts in hopes of producing the NIL clause (**a** result). Our notation for directed MBUILDing will therefore obey the convention that the question or problem be written as the first argument of MBUILD, and if a result is present, it is the "answer" to the question relative to that MBUILD.

We conclude this section with a few final examples:

I can't figure out what caused John to leave.




I can answer the question.




Notice here that we do not write $p \Leftrightarrow \text{MBUILD}$. Written this way, we are asserting that p has the ability or mechanism of thought, not that this mechanism can produce any results. Every normal human being can MBUILD.

"Can answer" is therefore signified by the presence of the result.


Can a newborn infant think?




Are you thinking about the question?



Can you answer the question?



I've concluded that I just can't think anymore!



8. Important Extras

8.1 Understand

We have not yet discussed the verb 'understand' largely because everything that we have been discussing heretofore can be considered to be an instance of 'understanding'. For example, the very process of assigning a conceptual structure to a linguistic string can be considered 'understanding' that string, as can the process of assigning a new conceptual structure a place or a tag in memory. Thus, in order to talk about the concepts that we assign to the verb 'understand' it was first necessary to have presented the preceding work.

The mental act of understanding is like the physical act of building a bicycle according to a set of instructions, or identifying a species of bird, using a bird watcher's manual. Corresponding to a set of instructions, or a bird guide, the memory has structures of beliefs, sets of related facts, bodies of knowledge. In English we have names for some of these structures, like "French" and "Physics" as in "I understand Physics." We can also use almost any nominal to mean a set of facts about that nominal when we say things like "I know people" and "I know skiing."

This information is used in understanding (i.e. identifying or interpreting) some input belief by the act MBUILD. MBUILD takes as its object the input, which serves as the focus of thought, guiding the direction of interpretation. MBUILD references the information stored in the IM, using the paradigms found there to analyze the input. The results are at least momentarily entertained (i.e. **CONCed**) and then placed either in the CP or the IM. We can see then that the same input may be analyzed differently when different information is found in the IM. We can also see that someone


may very likely not be aware of all the things that influenced his analysis of some input. Further, since IM serves as the store for those things that are on our mind (things like the current situation being experienced and remembered knowledge about objects in the situation) ,we see that MBUILD will be using, and hence be affected by, not only the bodies of knowledge that have been brought (i.e. **MTRANSed**) into IM, but by contextual features as well. For example, suppose several people recently have done me favors just prior to asking for a loan of money. In my IM then is a belief of the form "Someone helps someone else to increase their chances of borrowing money." If someone now- asks me if I know why John has been so friendly lately, I'm likely to say that John probably wants to borrow sane money. I have **become** suspicious, that **is**, **MBUILD** finds an ulterior motive type of belief in my IM, ready for application to new inputs.

The body-of-knowledge concept helps to explain the source of the different senses and levels of understanding. We can see that the nature of the understanding is constant, but that the information used changes with each sense. The levels of underst**anding** arise then **from** the levels of abstraction **we attach** to different bodies of **knowledge**. The kinds of rules used to understand French differ from the kinds of rules used to understand John's motives, but the underlying action of using information remains the **same**.

With the above description of the way we intend to treat understanding, the following representations should seem reasonable.

Underlying any use of "understand" there is always an action. When used in the atemporal sense, as in "I understand French", it is referring to a capability, as in "I speak French", and "I type 80 words a minute. "

That is, "I understand French" as normally used communicates the same thing as "I can understand French when I hear it." That is



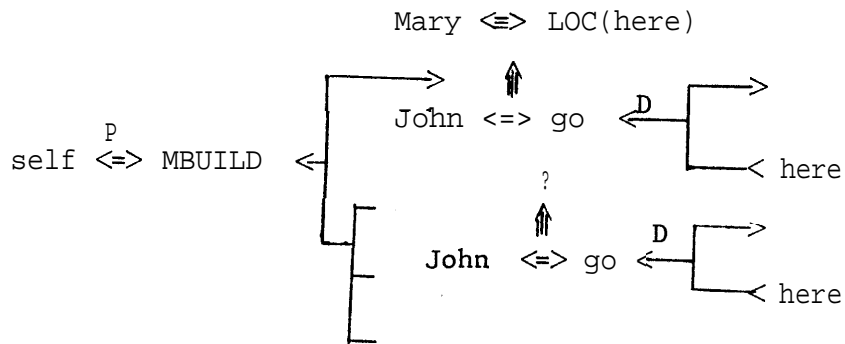
Associated with the fact that "I can understand French" is a stronger statement that "I know French." This is representable as

French \Leftrightarrow MLOC(LTM(self))

From this it is inferrable that, since one can always attempt to MBUILD, and since one knows that one has a knowledge of French by successfully MTRANSing it to CP, I can use this knowledge to understand and speak French.


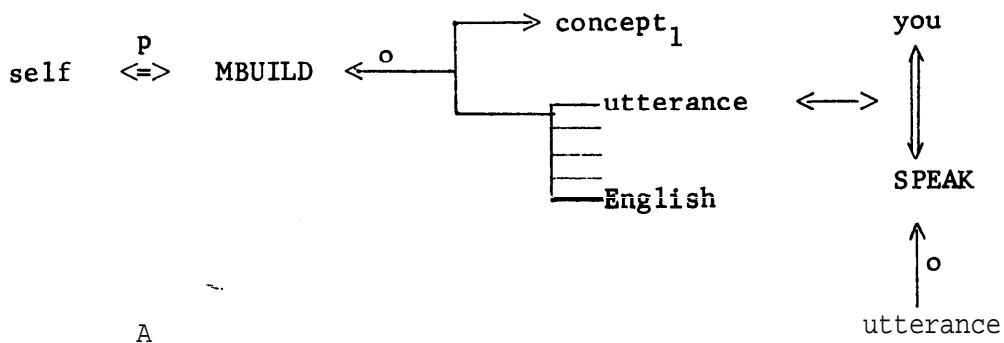
Understanding also involves MBUILD as the basic act in its non-capability sense.

"I understand why John left. Mary arrived."



The "?" is used here to represent a question, which is guiding the **MBUILDing**. In this case, the question is "Why did John leave?"

Since "understand" involves many different areas of knowledge, it can happen that it will be used twice in the same sentence in two different ways. This happens in "I understand what you said but I don't understand what you mean," which is:



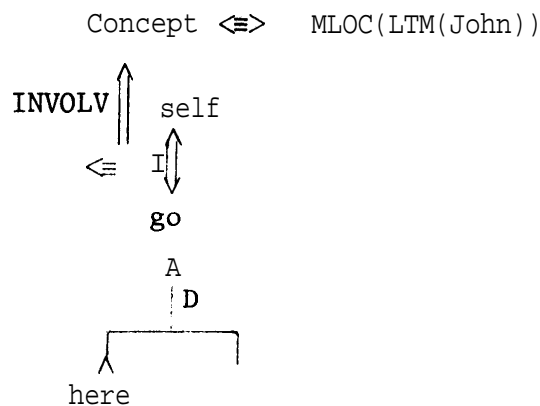
where "previous context" stands for a set of conceptualizations that have been built up during the conversation.

There seem to be two major concepts involved here -- one is a static item, the structure, while the other is a dynamic action, interpretation. A new element is said to be understood when it is **successfully interpreted**

according to some structure. The different senses of "understand" arise then from the different structures used by the interpreter, and the feeling of "levels of understanding" follows from the feeling of levels of abstraction associated with the various structures. Almost any noun can be used with "understand" or "know" in this body of knowledge sense. "I know London." "I understand people." "I know wine." Linguistically this is a generic use of the noun, but it is more than that since clearly "dogs" means one thing in "Dogs eat bones," and another in "I know dogs." There is ambiguity to the second sentence (usually not found when the stress pattern is included) as to whether it means "I know facts about dogs" or "I tell you who I know, I know dogs."

There is another static sense of "understand", that looks like "believe". This sense occurs often when the 3rd. person present tense is used with a question as a complement. For example:

"John understands why I left."



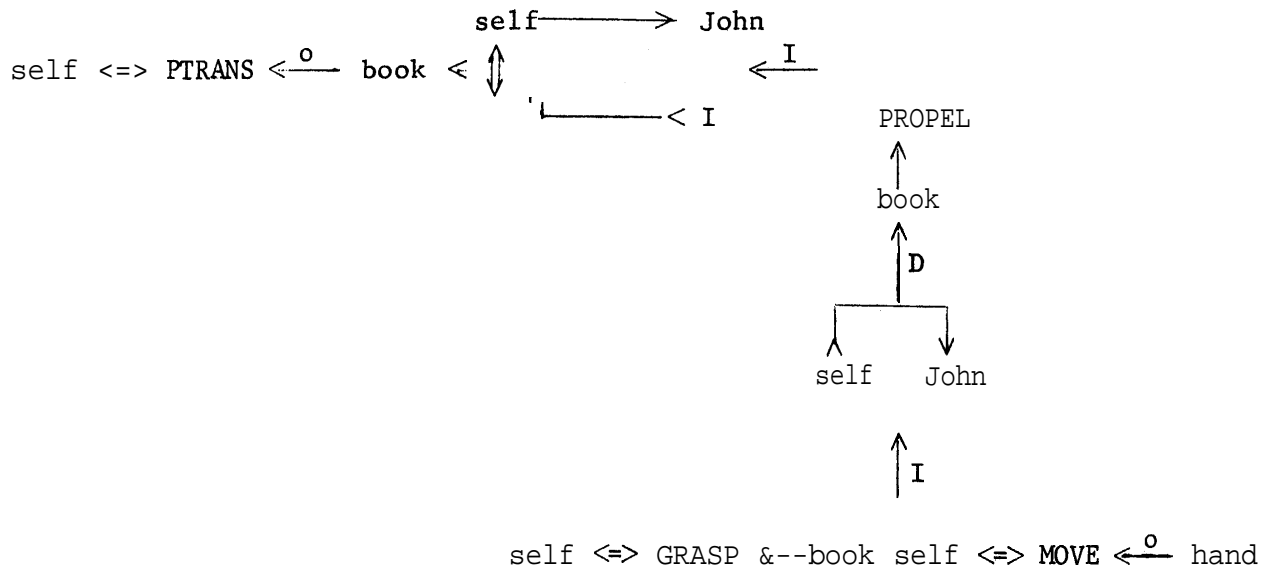
This is the result of the "understand" actions of MBUILD, like those performed in the earlier example "I understand why John left." Notice though that once MBUILD has produced a conceptualization, the holding of that conceptualization is knowing it. That is, whenever "understands" is used in the above way, it could be replaced by "knows". The nature of the question appearing in the complement may restrict, however, the substitution of

"understands" for "knows". Thus we could have used "understands" rather than "knows" in "John knows why I left" but not in "John knows when I left".


Sometimes this static sense occurs in sentences like "I understand French," and "I understand John." We are simply saying that we have a body of knowledge in our LTM. From this we can make inferences like "Someone can communicate to me in French," and "I can predict John's actions and motives," because we normally infer that if someone knows he has beliefs then he is capable of **MTRANSing** those beliefs to his IM.

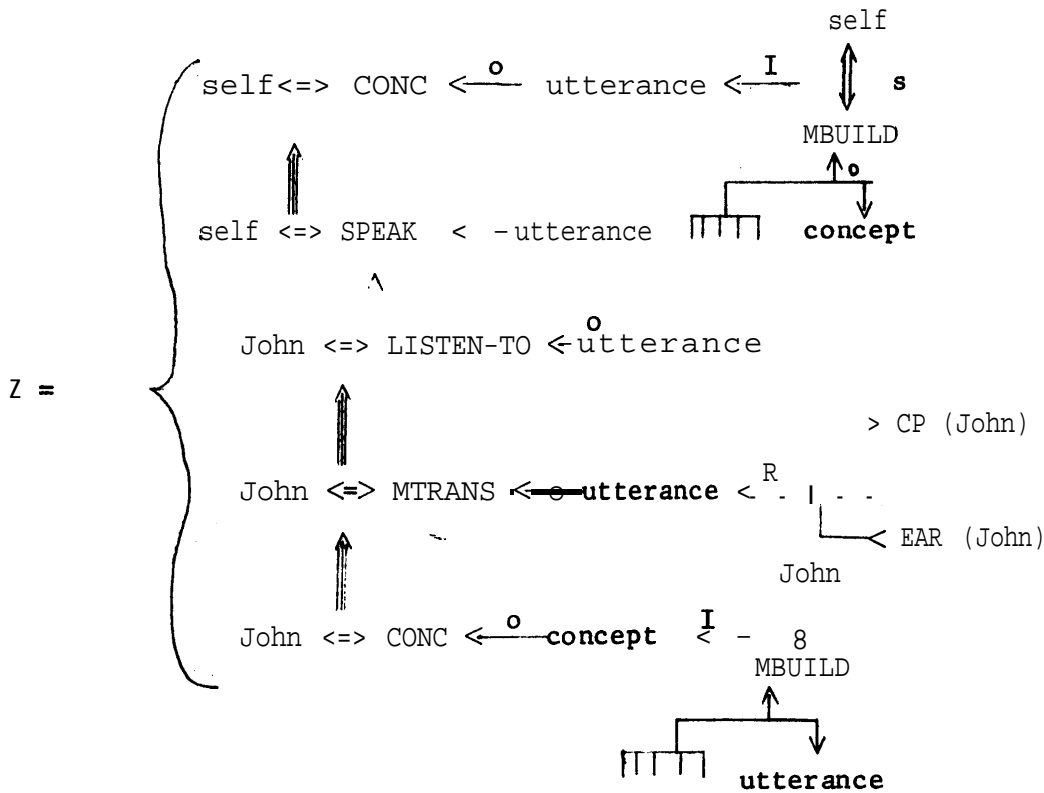
8.2 MTRANS vs PTRANS

With the apparatus provided by MBUILD and the idea of structions of knowledge we can strengthen the analogic relationship between MTRANS and PTRANS, as used in "communicate" and "give" respectively. PTRANS, it has been noted, can be broken up into constituent physical actions by use of the instrumental case:

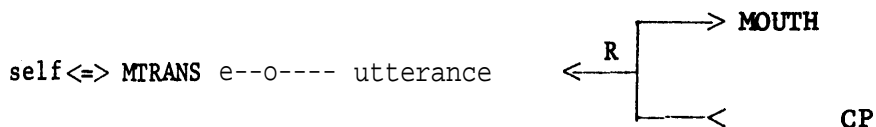


Similarly we can break MTRANS in "communicate" into its constituent mental and physical actions:





The asymmetry between speaking and hearing may indicate a need for something like:



but no such claim shall be made at the moment.

We can see that understand in "I understand what you mean" is one part of the total **communicative** action, **focussed** on the hearer's actions. When just MTRANS is used, we are emphasizing the point that it is the flow of information (which actually occurs in two **MTRANSes**) that is basic and constant to all forms of communication, just as PTRANS, the change of possession,


underlies all the actions of giving and taking. MTRANS and PTRANS are what the actor intends to do, by means of certain physical and mental actions.

8.3 Judging

Judging consists of assigning a value (relative to some norm) along a specific attribute dimension of an object or event. Sometimes the dimension is specified in an utterance -- "Please judge the monetary value of this diamond ." More commonly, the dimension must be inferred from the nature of the material being judged and the context in which the utterance is encountered -- "Please look at this diamond and tell me what you think of it ." In such a case a conceptual parser must apply its inference capacities in order to determine the appropriate dimension(s) before the conceptual content of the utterance may be determined.

Since the act of judging can be viewed as the attachment of a value to a conceptualization, we are able to represent it using the primitive ACT MBUILD . In the first case above we have:

"Judging the monetary value of a diamond"



We would likely represent the second utterance above analogously were it encountered in a pawnshop, while the dimension of judgement would be entirely different if the speaker were a prospective bride. In this case we might choose the dimension "beauty" for evaluation.

It is at first tempting to express acts of judging as attachment of

credibility measures (see section 5.1) to **conceptualizations**. Thus **judging** a painting would be to attach an appropriate credibility to a conceptualization such as:


painting \Leftrightarrow good

The element 'good' above can only reasonably refer to some abstract 'goodness' norm for paintings. But at least two problems arise if we try to explain judging a painting in terms of attaching a credibility link to this conceptualization:

1) if we attach a high credibility, we are really expressing the certainty of the judge that the painting is of 'normal goodness', not an opinion that it is a very good painting,


2) if we attach a low credibility, we are representing the fact that the judge strongly disagrees with the proposition that the painting is of 'normal goodness.' But we do not represent his judgement of the painting's quality (it may be good or bad).

The appropriate way to represent "Judging the painting to be poor" as a conceptual diagram is:




whereas "being certain that the painting is poor" would be graphed:

painting \longleftrightarrow goodness = low



This is not to say that judging never consists of attaching credibility values to conceptualizations. In fact the legal sense of judge presents precisely this case:

"To judge John in the case of Mary's murder"



Fillmore (1971) discusses a semantic representation for several verbs which he calls 'verbs of judging'. Most of these verbs -- e.g., accuse, criticize, praise -- express the communication of a judgement rather than the actual action of judging. (The fact that a judgement was made is indirectly indicated in Fillmore's discussion of the 'presuppositions' of these verbs).

Those verbs in Fillmore's list which do not necessarily refer to communicative events -- e.g., credit, blame -- express beliefs which are normally results of judgments.

8.4 Volition, Intention, and Their Relation to Actions

Often, what is being communicated is not simply the relating of an event, but also an indication that the event or the goal of the event was intended by the actor.

Volition and intention always imply the existence of some past, present or future action on the part of the actor whose volition or intention is involved. We do not, for instance, say things like "I intend for Neil to trip on the volleyball court tonight, but I won't do anything to cause it." This is a simple but important observation, for it tells us we must always account for such an action in the conceptual representation of verbs of volition and intention. Also, while it may seem on the surface that we intend to cause a state to exist ("I plan for our barn to be orange."), conceptually we intend to perform an action which will bring about that state (i.e., the painting of the barn, the hiring of a barn painter, etc.).

The following is a representative list of conceptual structures related to volition and intention which we feel can be characterized in terms of CONC:

"I did it intentionally."

"I did it consciously."

"I was aware of doing it."

"I did it accidentally."

"~~I did it~~, but I didn't want to."

"I intended to do it, and I did."

"I intended that X occur, and it did."

"I intended to do it, but didn't."

"I intended that X occur, but it didn't."

"I intend to do it."

"I plan to do it."

"I am trying to do it."

"I will try to do it."

'(I tried to to it, but did not succeed."

We will require three related CD graph **structures** to account for these sample constructions: Simple volition, intention: past and present tense usage and intention: future tense usage .

VOLITION


The term "volition" ("intentionality", "willfulness", etc.) has a relatively precise -conceptual underpinning: it refers to an actor's CONCEPTUALIZATION of his potential action before he performs that action, Contrast the sentences:

"I dropped the cat."


"I dropped the cat accidentally."

"I dropped the cat intentionally."

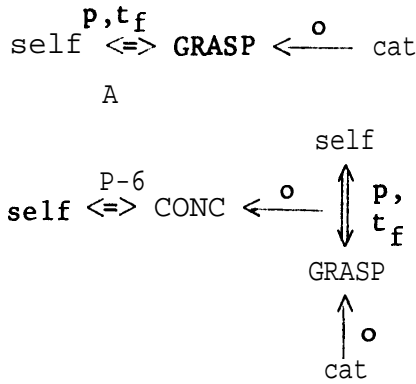
The first makes no statement about the presence or absence of volition, and is therefore represented as simply




The second makes explicit that it was not an intentional act, and is hence written:




and the third makes explicit the intentionality of the act:



Notice that although the physical cause of the dropping is not stated in any of these, it is always implied. When volition is not involved, we would ascribe the dropping to some "external" cause: "The cat squirmed loose", or "I didn't have a firm **grip**." But what is the cause when volition is involved? "Clearly" it is the firing of a group of neurons in the brain which causes impulses to travel down the arm to the fingers, causing muscle cells to fire. Admittedly, the original cause of free will is a question better left to the philosophers. But, as semanticists, we cannot resist our natural urge to argue that the firing of the neurons in the brain was the original cause. This is especially attractive, since such neural activity is precisely one of the micro-mental processes for which **CONCeptualize** was postulated. Viewed this way, a **CONC** in the mental domain can cause actions in the physical domain, so that the volitional cat-dropping incident can be expressed as:



This scheme

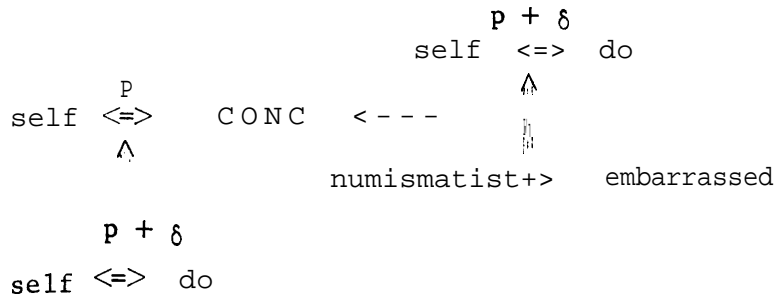


turns out to be a very general CD template, recurring as the central theme of all volitional acts.

INTENTION(past and present tense usage)

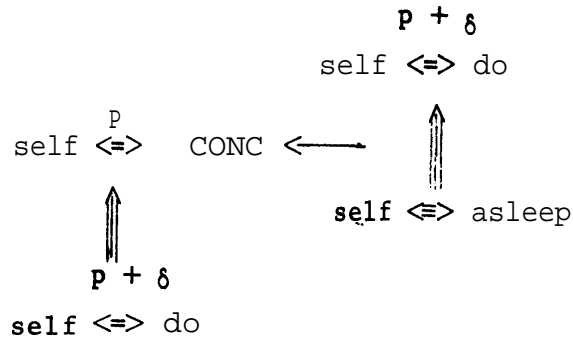
There are forms of volitional activity where the **CONC**eptualizing of more than just the future action is the cause of that action. In these forms, rather than **CONC**eptualizing just the action itself, what is **CONC**eptualized is a causality relation involving the action as cause. In other words, we become aware that performing an action will cause some (desired) goal, and 'it is the awareness of this causal relation which causes us to act. In English, this paradigm is often represented by "intend" or "try". Some examples are:

"I intended to embarrass the numismatist."



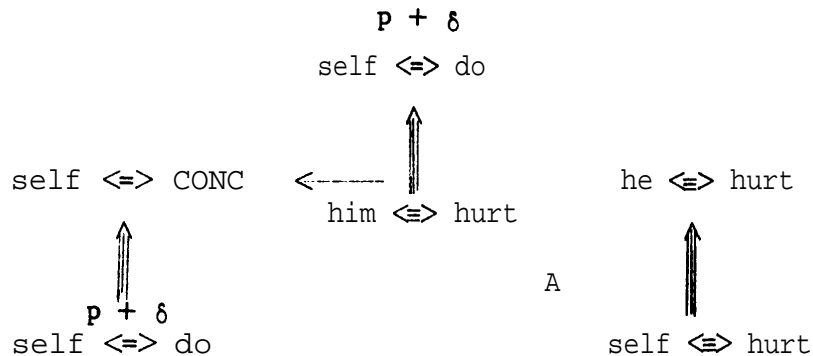
and

"I'm trying to fall asleep."



There is the question of desirability of the goal (result of the **CONCEptualized** causal) in this form of intend. Are we always to infer that whenever one performs an intentional action (which he **CONCEptualizes** will lead to some **result**), the result is desirable to him? If this is the case, we would be **justified** in omitting the explicit mention of the result's desirability, leaving this as an inference to be made whenever this structure occurs. There would appear to be cases where the result is intentional, but not desirable:

"Although it hurt me to do it, I intended that he be hurt."




However, by writing such things, we are concealing some information. We view people as always seeking goals which are **CONCEPTUALIZED** at the time as leading to the most ultimate pleasure (least displeasure). Viewed with such stalwart hedonism, it must be concluded that a volitional ACT driven by the cognizance of what the act can cause can only mean that the result is desired. Therefore, we may always infer the desirability of the result. Notice that in cases such as the above, we are left with an apparent contradiction that the result is both desirable and undesirable. But this is precisely what we want. It indicates that the desirability of the result must be assessable on either of 2 levels, and that there is really no contradiction at all. The presence of such a "contradiction" provides a meaningful "level-disambiguation" task to the understanding program using this theory.

INTENTION (future usage)

So far, we have restricted the discussion to past and present usages of volition and intention. In these forms, an action is either performed or attempted. In contrast, future use of intention connotes only that an action is planned in the future and makes no prediction about its success or even its attempt. In the future use, we make use of the propositional nature of a potential action, predicating that it will be true at some future time:

"I intend to leave for McLean tomorrow."





Notice that a) there is no causality involved, and b) what is **CONCEPTUALIZED** is not the action itself, but the fact that the action, viewed as a proposition, will be true. This is quite different from volitional actions where the pre-**CONCEPTUALIZATION** of an action itself is responsible for the action. Conversely, ACTS

of volition are not performed by **CONCEPTUALIZING** that their value as a proposition will be true, but rather by **CONCEPTUALIZING** the action itself.

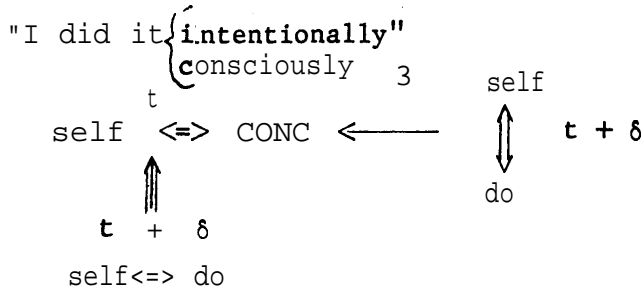
This future construction may occur in the guise of a syntactic past tense construction. This happens when we say "I intended X, but never really got around to doing or causing it." Seemingly similar to the past use of intend, where an action was attempted or performed (but no result obtained), this use of intend is conceptually quite different, since no action was even attempted. An example:

- "I intended to embarrass the numismatist, but
 ...
 a) I never got around to it."
 b) I decided not to."




EXAMPLES

We conclude with semantic renderings of the sample sentences presented at the beginning of this section.

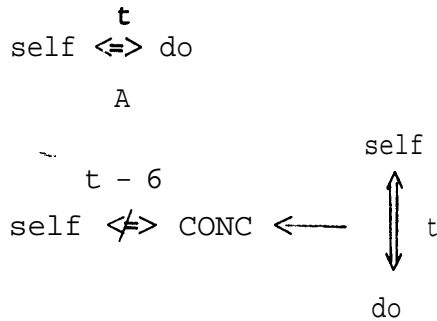


"I was aware of doing it."

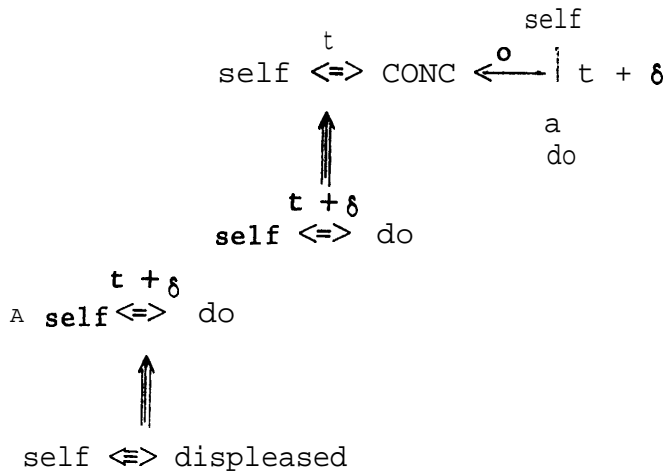


(i.e., awareness is not volition: it is the perceiving of an action, rather than the will to act.)

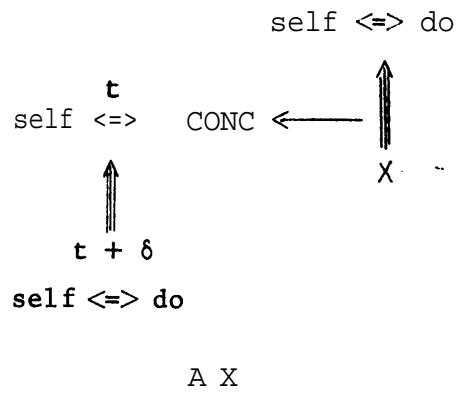
"I did it accidentally."



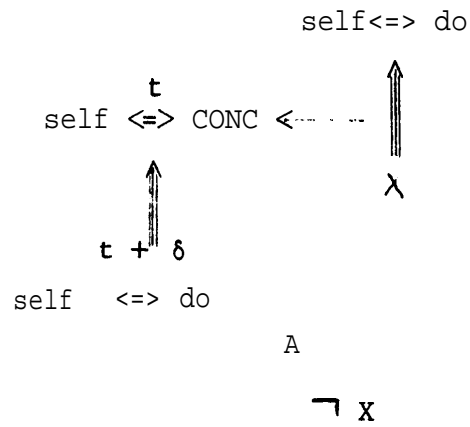
"I did it intentionally, but I didn't want to."




"I intended that X happen, and it did."




"I intended that X happen, but it didn't."




"I $\left\{ \begin{matrix} \text{intend} \\ \text{plan} \end{matrix} \right\}$ to do something."



"I will try to do something."




" I am trying to bring it about."




('Try' = intend, but emphasizes the action more. 'Try' in the past always indicates the action occurred, 'intend' in the past, only that the action was conceptualized.)


"I intended to do it, but I didn't."



"I tried to go home."



(A normal inference when try is used in the past tense is that the paradigm:



did not lead to the result.

9. Conclusion

It is perhaps reasonable at this point to attempt to put the discussions that have been presented here into the correct perspective. We make no claim as to the ultimate correctness of the material or particular analyses presented here. Rather, we have been trying to present a point of view rather than a set of individual verb-analyses. Basically, we have been trying to establish that looking at language conceptually affords some unique advantages. With respect to linguistics, the main advantages have to do with the replacement of the problems of whether something is grammatical or permissible to say by the problems of how to interpret any spoken utterance, and how to judge whether it makes sense to think a given thought. This leads one naturally to the problems of establishing a syntax of conceptual items that can aid enormously the problem of understanding what goes on in language. The problem of basic ACT primitives follows naturally from the conceptual syntax and the need to express meanings that are conceptually identical in one and only one way. This problem of the establishment of such primitives has been beginning to make itself clear to some researchers who have in the past taken a more traditional (i.e., transformational) approach to the problem. In particular, Miller (1971) and Fillmore (1971) have been considering the problem of what underlies verbs. While we find **this work** encouraging we must point out that it is entirely necessary to have a formal conceptual syntax before attempting the problem. Once it is clear what the nature of an ACT must be, it becomes much clearer how to go about establishing such an ACT primitive. Once it is realized that only whole conceptualizations can cause events then the tendency to make the transitive verbs 'hurt', 'fly', 'move', 'grow', 'bother', etc. into

primitive **ACTs** is greatly diminished.

With respect to the computer science approach to language, once a conceptual syntax is created that handles a given meaning in only one way, the problem of inference becomes clearer. If identical meanings with vastly different surface forms can be mapped into **only** one conceptual structure, a significant amount of the inference problem is eliminated. Furthermore, it then becomes possible to structure the memory that will operate on top of the language analysis programs in terms of the conceptual syntax.

This paper has been an attempt to arrive at a set of ACT primitives that will facilitate the solution to the above problems. We have intended to demonstrate how much can be done by looking at language in this manner and attempting this problem. Certainly there are a great many remaining problems of which we are aware. Consider, for example, the analysis presented for the verb "see". We have said that "see" is conceptually: CONC X by **MTRANS** X by LOOK-AT X. In actuality, the X in each case is quite different. As object of LOOK-AT, X is a real life physical object. The thing that is **MTRANSed** however is the perceptual image of X. The object of CONC is yet again different in the Xs? It must be the case that an ACT is present that is changing the form of the X. Shouldn't we posit an ACT to account for this? If we did we would be dealing with the world as it actually exists rather than as people talk about it. This is something that we don't want to do since we are discussing what people talk about and not what is "really" going on. If we allowed ourselves the luxury of dealing with only what is really going on we would be doomed to failure, partially because no one really knows what is really going on in the brain, and partially because even if we did we **would** be forced to deal with neurons

and brain impulses, a most undesirable situation. So, in lieu of representing what is actually happening, we must deal with what is happening as people talk about or perceive it. This leaves us with such things as an MTRANS that can magically change its object. A similar situation exists with PTRANS. The analysis of physical events in terms of what actual actions have occurred leaves out the ACT **PTRANS** simply because it is not something that can actually be observed in the world. Rather, PTRANS is an action that people talk about and is real in their mental world without having any overt physical reality. That is, PTRANS is a kind of **culturally** defined concept that an observer from a culture that does not have the concept of **possession** would not see. Thus, it is not an ACT in the sense that it is something that actually occurs in the world. But it is something that occurs in our mental perception of the real world and it is this that we talk about in language.

Thus, we have been trying to set down all the real primitive mental actions that people talk about. The physical world also has a set of real actions that people talk about, only some of these (like PTRANS) are actually purely mental.

.



REFERENCES

- 1] Becker, J. "An Information-Processing Model of Intermediate Level Cognition", A.I. Memo No. AIM-119, Computer Science Department, Stanford University, Stanford, California, May 1970.
- 2] Fillmore, C., "The Case for Case" in Bach and Harms (eds.) Universals in Linguistic Theory, Holt, Rinehart and Winston, New York, 1968.
- 3] Fillmore, C., "Verbs of Judging: An Exercise in Semantic Description" in Fillmore and Langendoen (eds.) Studies in Linguistic Semantics, Holt, Rinehart and Winston, New York, 1971.
- 4] Katz, J. and Fodor, J., "The Structure of a Semantic Theory", Language 1963, Vol. 39.
- 5] Miller, G., "English Verbs of Motion: A Semantic Analysis Based on Incomplete Definitions", Mimeographed, May 1971.
- 63 Postal, P., "On the Surface Verb "Remind"", Linguistic Inquiry, 1970, Vol. 1, No. 1.
- 7] Price, H., Belief, George Allen and Unwin Ltd., 1969.
- 8] Quillian, R., "Semantic Memory", Bolt, Beranek and Newman, 1966.
- 9] Schank, R., "A Conceptual Dependency Representation for a Computer-Oriented Semantics", A.I. Memo No. AI-83, Computer Science Department, Stanford University, Stanford, California, March 1969.
- 10] Schank, R., and Tesler, L., "A Conceptual Parser for Natural Language", Proceedings of the International Joint Conference on Artificial Intelligence, May 1969.
- 11] Schank, R., Tesler, L., and Weber, S., "Spinoza II: Conceptual Case-Based Natural Language Analysis", A.I. Memo No. AIM-109, Computer Science Department, Stanford University, Stanford, California, January 1970.
- 12] Schank, R., "Semantics" in Conceptual Analysis", A.I. Memo No. AIM-122, Computer Science Department, Stanford University, Stanford, California, May 1970.
- 13] Schank, R., "Intention, Memory and Computer Understanding", A.I. Memo No. AIM-140, Computer Science Department, Stanford University, Stanford, California, January 1971.
- 14] Schank, R., "Conceptual Dependency: A Theory of Natural Language Understanding", Cognitive Psychology, (in press), (1972).

REFERENCES (Continued)

- 15] Winograd, T., "Procedures as Representation for Data in a Computer Program for Understanding Natural Language", **TR-84**, M.I.T. Project MAC, 1971.
- 163 Yinger, J., Toward a Field Theory of Behavior, McGraw-Hill Book Company, 1965.