

STAN-CS-72-290

AD 746147

ADVERBS AND BELIEF

BY

ROGER C. SCHANK

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

JUNE 1972

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U S Department of Commerce
Springfield VA 22151

COMPUTER SCIENCE DEPARTMENT
School of Humanities and Sciences
STANFORD UNIVERSITY

ADDC
RECEIVED
AUG 7 1972
RECEIVED



This document has been approved
for public release and sale; its
distribution is unlimited.

ADVERBS AND BELIEF

BY

Roger C. Schank

ABSTRACT: The treatment of a certain class of adverbs in conceptual representation is given. Certain adverbs are shown to be representative of complex belief structures. These adverbs serve as pointers that explain where the sentence that they modify belongs in a belief structure.

Department of Computer Science
Committee on Linguistics

This research is supported by Grant PHS MH 06645-11 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 author 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 National Technical Information Service. Springfield, Virginia 22151.

Adverbs and Belief
Roger C. Schank

1. Introduction - The prevailing viewpoint in current linguistic theory, whether standard transformational theory or generative semantics, has been to see language in terms of a device for accepting the sentences of a language and assigning a structure to those sentences with regard to their meaning. This point of view has been applied within what has been called a competence theory.

It is, of course, possible to look at language from other viewpoints. We can, for example, consider language to be a device for transmitting conceptual information between people. Linguistic analysis, then, would be in terms of providing the formal representations for conceptual information and the explicit rules for both the decoding of linguistic strings into these representations and the encoding of the information represented conceptually into linguistic strings. Such an approach would not be concerned with the accepting of sentences then, but rather with their interpretation and production. This kind of theory has been called a performance theory by generative linguists. Such a label brings to mind things like inattention and false starts (as stated by Chomsky [1965] in his discussion

of performance theory) which is most certainly not what such a theory seeks to explain. We choose, then, to forego the competence-performance distinction, and to refer here instead to a theory of language understanding.

The point of this paper is that such a theory of language understanding makes explicit certain implicit relationships present in language that have to date not been handled by generative theories.

2. The Conceptual Level - The particular topic with which we shall concern ourselves here is a certain class of adverbs. Primarily, generative linguists have considered in their discussion of adverbs the distinction between those that modify the verb and those that are sentence modifiers. (For example, see Lakoff [1970a] and [1970b].) While generative linguists are quite concerned with the place of adverbs in a semantic structure that reflects meaning, they rarely concern themselves with what the adverbs themselves mean.

To elaborate upon this, it is necessary to introduce a deeper level of linguistic description than is in common use, which we shall call the conceptual. Roughly, then, we shall say that there are three levels of description: the syntactic, the semantic, and the conceptual. Consider sentence (1):

- (1) John threw a hammer at Bill vengefully.

Discussion about this sentence on the syntactic level might center around whether "vengefully" should be placed before the verb or after the whole sentence in order to be grammatical. On the semantic level, we might concern ourselves with the question of "vengefully" as a predicate modifier or a sentence modifier. On the conceptual level, however, we are concerned with the meaning of "vengefully". That is, paraphrased in some other terms, what representation of "vengefully" would make clear the conceptual information that is imparted by this word.

Clearly, then, the conceptual level is not a level of purely linguistic description. That is, we do not seek, in our representation at the conceptual level, to represent the relationship between linguistic entities but rather the relationship between conceptual entities. That is, items which may not appear at all in a given sentence can certainly appear in the conceptual representation underlying that sentence. As a simple example of this consider sentence (2):

(2) John bought a book from Mary.

The conceptual representation underlying (2) must have in it the information that "John gave some money to Mary which caused Mary to give a book to John". Now, of course, the first obvious difference between this and a semantic description is that the idea of "money" is present in the conceptual representation whereas it is not in the surface sentence. It is of course true that in a possible inter-

pretation of this sentence, "money" may not have been present at all. But, it is the responsibility of the conceptual level to represent the most likely interpretation of a given sentence within a context. Given the highly artificial nature of linguistics papers, it is pointless to debate about what a given sentence might mean. However, it is extremely important for any mechanism that is intended to operate in context (as is any language understanding theory) that it be able to come up with any assumed implicit information that is not explicitly stated in the complete contextual situation. That is, the conceptual level serves as a representation for the implicit and explicit information underlying a surface string.

In addition, we claim that any two sentences that are said to have the same meaning may have different semantic representations (that is, they may use different words or sentence forms) but they must have identical conceptual representations. Therefore, it is necessary to establish a set of primitive concepts into which semantic structures that have the same meaning can be mapped. For sentence (2) we use the conceptual primitive action TRANS for the "giving" action. Every action (ACT) requires three or four out of five possible conceptual cases (Actor [A], Objective [O], Instrumental [I], Recipient [R], or Directive [D]). (We shall not go into the requirements of the conceptual level here as this is fully explained in Schank [in press].) In

addition, 'cause' is treated as a relationship, not an ACT, and represented by a dependency arrow $\widehat{(\uparrow\uparrow\uparrow)}$ between causer and caused. The basic composite unit of the conceptual level is called a conceptualization (which is denoted by a C and a number followed by an ACT and a set of conceptual cases). The relationship between the conceptualizations is given in the first part of the conceptual diagram. The actual elements of the conceptualization are given in the second part.

A primitive ACT is written in capital letters in first position in a conceptualization. If the first item is not in capital letters, it is a state, not an ACT. Thus, the conceptual representation for (2) is:

		A	O	I	R
$\widehat{C1}$	C1: TRANS	(John, money,			, Mary)
$\uparrow\uparrow\uparrow$					
C2	C2: TRANS	(Mary, book ,			, John)

Since the point of this paper is to discuss adverbs, we shall not go into the nuances of conceptual representation. (The notation used above and throughout this paper is considerably different from that used in previous papers describing this work. This is done for the sake of readability.) Similarly, we shall not discuss the particular primitive ACTs used but rather refer the reader to Schank et. al. [1972] for a discussion of the sixteen primitive ACTs that are used at the conceptual level.

It will be necessary, however, to discuss the basic idea of what a conceptualization is and therefore what qualifies as an ACT. Consider sentence (3):

(3) John hurt Mary.

We claim that in order for something to qualify as an ACT it must be invariant regardless of the sentence in which it was contained. Thus, "hurt" in (3) is not an ACT because what John actually did to hurt Mary is variable. That is, John may have kicked Mary or insulted her mother or whatever. What we do know is that this variable (i.e. unstated) action of John's resulted in a given state, "hurt", (although that state is actually ambiguous between mental and physical hurt). Thus our first actor-action-object conceptualization underlying (3) must have a variable ACT in it (which we call DO). Thus C1 is "DO John". The second conceptualization underlying (3) is a state relationship between Mary and hurt (C2). The relationship between C1 and C2 then is causality, that is, C1 caused C2. We write the conceptualization underlying (3) as:

		A
C1	C1: DO	John
		
C2	C2: hurt	Mary

Notice that here we are treating "cause" as a relation

rather than as an action as is traditionally done. The reason for this is that if we used "John caused Mary to be hurt" we would be missing the important idea that John did something that was unstated. It is this unstated action that caused the resultant state. It is also more obvious now what to do with any instrumental phrase that might occur. For example in the by-phrase (4):

(4) John hurt Mary by kicking her.

"kick" replaces the DO in the above conceptual diagram.

3. Adverbs - Now we can return to sentence (1). On the conceptual level, we consider the underlying ACT for "throw" to be PROPEL, meaning "apply a force to". The means by which this PROPEL-ing is accomplished is considered the instrumental conceptualization of PROPEL. (Conceptually, instruments can only be complete conceptualizations, never a single object. When the action that was used on that object is unstated it is usually possible to infer it.) The instrumental actions for PROPEL in the case of "throw" are MOVE (where the object is a hand containing the hammer (written hand CONT hammer) and UNGRASP (where the object is the hammer).

So, without the word "vengefully", sentence (1) would have the conceptual diagram:

	A	O	D	I
Cl:	PROPEL	John	hammer	Bill Cla + Clb
Cla:	MOVE	John	hand CONT hammer	Bill
Clb:	UNGRASP	John	hammer	

(We read this as Cl by means of Cla followed by Clb.)

Now the interesting question is, how does "vengefully" affect this structure? If we simply modified the main ACTS involved (PROPEL and MOVE) we would explain nothing since we can consider something like "MOVE vengefully" to not be primitive at the conceptual level. That is, this would have to be broken down in order to explain what it means (the task of the conceptual level in the first place). The only possible modifiers of primitive conceptual actions are those that actually refer to aspects of those actions. Consider "MOVE". The primitive action MOVE is used whenever a body part 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" there are, 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, for the second sense of (1) can vengefully 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 which we shall mention at

the end of this discussion), we have to find some other place for it.

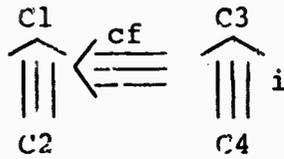
It is important to realize that "vengefully" is simply another form of "revenge". In order to deal with a meaning analysis of the concepts of a sentence containing "vengefully", it is necessary to deal with the meaning of "revenge". "Revenge" is not a simple word by any means. The reason for this is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

We define belief 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 "X is one who is likely to do Y", or "X is bad", and so on. Since language is a means of expressing beliefs, it is only right that in doing a conceptual analysis of a linguistic expression that we explicitly state the beliefs that are implicit in that expression. (We should point out here that these beliefs are only a small part of what are commonly labeled "beliefs". Beliefs of the order of "I believe that John hit Mary" shall not be discussed here.)

Computer programs have been written (see in particular Colby et.al. [1971] and Abelson [1965]) that use beliefs to simulate human thought behavior which have tried to avoid

the problem of analyzing natural language expressions. We are claiming in adding the notion of belief to natural language analysis that analysis of linguistic input is done with some purpose, and in most simulations of human behavior on a computer the language analysis must be done with regard for the purpose for which the hearer is doing the hearing. This is not an odd statement to psychologists doing such modeling of belief (certainly the two authors cited above would readily agree with it) but rather is one that is traditionally odd to linguists. Psychologists doing computer modeling of human behavior have avoided dealing with natural language because of the difficulty of doing that rather than because they thought that they shouldn't. Linguists' avoidance of the psychological expressions inherent in natural language has been caused more by an attempt to analyze language by itself in some sterile environment. Unfortunately, natural language exists in people's heads, rather than in a vacuum so it is not unreasonable to be concerned with analyzing natural language utterances with respect to their global content. Thus, it is not unreasonable to make explicit the psychological statements that are being made in sentences.

With respect to sentence (1), "vengeance" can be said to be reflective of the following belief-conceptual structure.



A

- C1: DO one₁
- C2: hurt one₂
- C3: DO one₂
- C4: hurt one₁

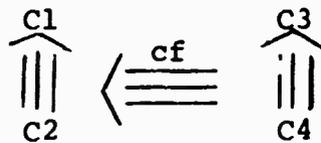
The above structure is to be read as: "the causal relationship between C1 and C2 could (c) cause in the future (f) the relationship C3 is intended (i) to cause C4'. What we are saying, then, is that if person 1 (one₁) causes person 2 (one₂) harm, this could cause person 2 to do something that is intended to harm person 1 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 conditions. The structure given is simply that elicited by the word "revenge". The words "vengeance" and "vengefully" call this structure as well. Moreover, when the word "vengefully" is present, the conceptualization underlying the sentence that "vengefully" modified can be placed in the C3 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, avenge:

Each word calls the following belief:



		A
C1:	DO	one ₁
C2	hurt	one ₂
C3	DO	one ₂
C4	hurt	one ₁

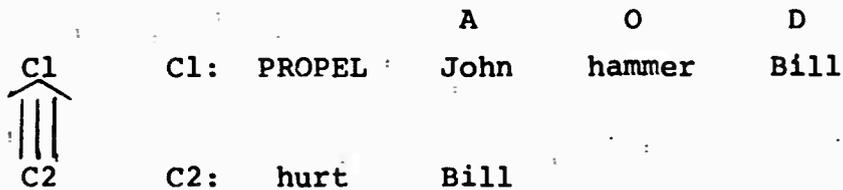
The conceptualization that is modified by the word under discussion in the sentence is to be placed as C3 in the above belief.

Notice that the above entry reads like a command to a memory system rather than your usual dictionary entry. That is precisely what it is. Since the dictionary we are referring to is to be used for analysis into conceptual structures it is often the case that the entries turn out to be commands to the system to do something with respect to the creation of such structures.

We have not quite finished with "vengefully" yet, however. It is a common inference on the part of hearers

of sentence (1) that John probably hit the man hard. It is interesting to examine where this inference comes from.

As we have stated, it would be incorrect to simply claim that "vengefully" means "hard". Rather, if it does mean hard it is because of the meaning of "vengefully". We have shown that "vengefully" means in part that the sentence under consideration fits into a structure that indicates that this conceptualization had as its intention the "hurting" of the object of "hit" in the sentence. Thus we have:



Since we know that the intention of the propulsion of the hammer is the hurting of Bill, we can make an inference from our knowledge of the world about the strength of the propulsion. Bear in mind that what we can make is an inference and in that sense it is only a probable statement and by no means must it be valid.

The reasoning that people do is something like this: If he wanted to hurt the man, then since the harder you throw something the harder it hurts, he probably threw the hammer hard. This can be written more formally as: (1) Physical hurt requires force. (2) $F = MA$. (3) $F = M_{OBJ} A_{Action}$.

(4) ∴ add modifier of "speed" to "MOVE". Thus the inference of "speed" can be added as a modifier to "MOVE" conceptually.

This, however, is only an inference about this conceptual structure and is quite a different thing than the previous statements about the meaning of "vengefully". The meaning of "vengefully" conceptually is what we have stated it to be (above) under all conditions. But an inference can be, and often is, wrong.

One reason is because of the ambiguity of the belief structure used above. "Hurt" is ambiguous in our conceptual system. We have been referring here to one type of hurt, the physical type, written HURT_{PHYS}. Another type of hurt is HURT_{MENT}. Mental hurt is not usually caused by force, so we would never apply rule (1) above if we knew we had an instance of HURT_{MENT}. Thus, we are making a guess that we have physical hurt due to the physical nature of the props (i.e., "hammer") of this sentence. That is, context influences our selection of senses of a word. Since we have somewhat artificially created a sentence without a situation, we can only guess what might be correct. In a known context, if we have hurt_{PHYS} the inference is likely to be valid. If we have hurt_{MENT} it will never be made.

Thus, certain adverbs are actually indicators that the conceptual structure underlying the verb of the sentence that the adverb modifies fit into a certain spot in a

complex belief structure. Often inferences can be made about those structures to yield some probable information about the initial conceptualization.

Let us examine some other adverbs that this statement is true for and what particular structures are represented by those adverbs. Consider sentence (5):

- (5) Mercifully, the King only banished the Knight for killing his favorite horse.

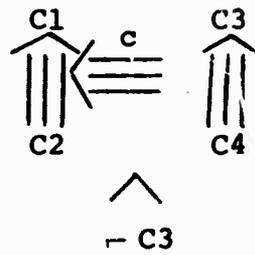
Here, "mercifully" references the same belief that was the object of our discussion of sentence (1). That is, we have something that is paraphrasable by: The Knight did something to hurt the King which could have led to the King hurting the Knight a great deal, but the King only hurt him a little. Or in other words, the punishment was less than might be expected.

Contrast sentence (5) with sentence (6):

- (6) Mercifully, the King let the Knight go unharmed.

Here, the punishment that we might have expected did not exist at all. Notice that the "only" is necessary in (5) in order to use "mercifully". That is, one hardly seems merciful if one is hurting someone. It is only "merciful" by comparison to some expected greater hurt.

So the belief for "merciful" is as follows:



(where \wedge means "and")

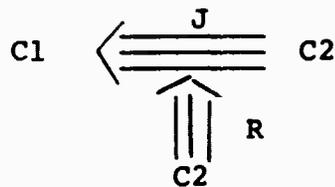
(where \neg means "not")

		A
C1:	DO	one ₁
C2:	hurt	one ₂
C3:	DO	one ₂
C4:	hurt	one ₁

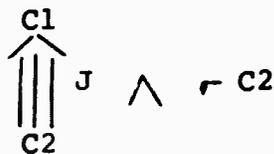
What we are expressing here is a basic mode of human thought. The structure is of the form: even though "X" justifies "Y", the actor did not do "Y". Conceptual structures of this form lead us to modify our initial conception of the correct conceptual structure for vengeance. It would be more accurate to have for vengeance: since "X" justifies "Y" it is true the actor did "Y". That is, we must link up a conceptualization with the reason for its occurrence. We thus define a justification as being of the form of "given C1, then we might reasonably expect C2", or:



This new link $\Uparrow J$ is called a justification. A justification is only true with respect to a given system. That is, when we use $\Uparrow J$ we are saying that the person about whom we are talking probably feels that C1 justifies C2. Such justifications are simply a certain type of belief. Justifications are used in conjunction with reasons. We are now stating that a justification can be a reason which we denote $\Uparrow R$ for an action. That is, we can have the following form:



Here, we are saying that since C1 justifies C2 in the actor's belief system, this was the reason (R) for his doing C2. Thus, words like "vengeance" can be defined in terms of this construction. "Mercy" on the other hand is of the form:

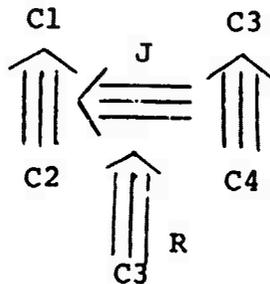


(where $\neg C2$ is interpreted to mean "C2 did not occur" [note that " $\wedge \neg$ " is "but" in English]).

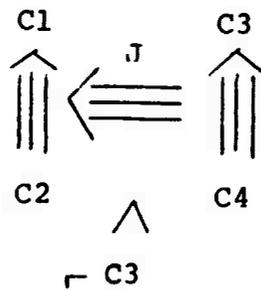
Now in fact we have just defined "reciprocity" rather than "vengeance". "Vengeance" implies that C1 and C2 were

bad things (i.e., they caused "hurt"). Thus, we have:

		A
	C1: DO	one ₁
FOR	C2: hurt	one ₂
	C3: DO	one ₂
	C4: hurt	one ₁

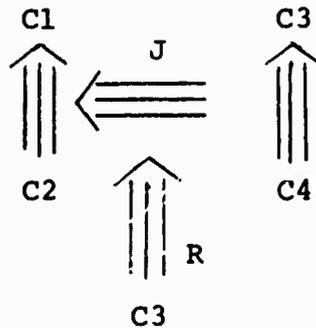


vengeance
revenge
get back at
avenge
reciprocate (bad)

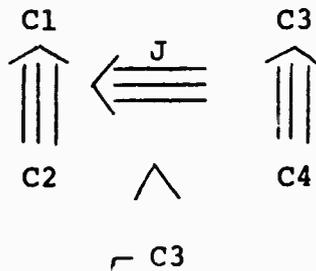


mercy
not get back at
kindly not do
generously not do

		A
	C1: DO	one ₁
FOR	C2: pleased	cne ₂
	C3: DO	one ₂
	C4: pleased	one ₁



return a favor
reciprocate (good)



be ungrateful
mooch

Another significant type of behavior is represented by the class of adverbs that refer to the reasonableness of a given action. For example consider:

- (7) Stupidly, I let him go.
- (8) I foolishly paid him too much money.
- (9) John thoughtlessly told Mary to go kill herself.
- (10) Fred hit John needlessly.

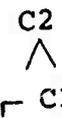
All of these adverbs refer to the reason and intended or actual effect of an action. Let us consider (10) first. How can an action be needless? An action is only needless if one of two possibilities hold. If the intended effect of the action did not occur (i.e. the goal of the action was not accomplished) then the action was needless. Or, if the goal of the action was accomplished by some other event then the action was needless.

In order to explain how to do this conceptually, it will be necessary to first introduce our method of dealing with intention. (10) implies that an intended goal exists and it is thus necessary to explain it. (The [i] used above for intention was merely a shorthand device.)

The basic ACT of thinking in conceptual dependency is CONCEPTUALIZE (CONC). We use CONC to indicate that a thought is being consciously processed at the time or "thought-about". In order to intend to do something it is thus necessary to CONC the thought of doing it first. Furthermore, most intended actions have intended goals. It is thus necessary to think about the causal effect of the CONC-ed action. We thus treat "intend" as the doing of an action that was preceded by the CONC-ing of that action and its effect. Thus "I intended to do C2 which would have effect C3" is:

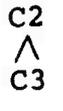
		A	O
C1: CONC		self	
C2: DO		self	
C3: state		object	

We can now get back to our example. Sentence (10) in one sense then (goal not accomplished) is thus:

		A	O
	C1: CONC	Fred	
C2	C2: HIT	Fred	John
	C3: unstated		

That is, the thought about effect did not occur.

In the second sense, we have (goal already accomplished):

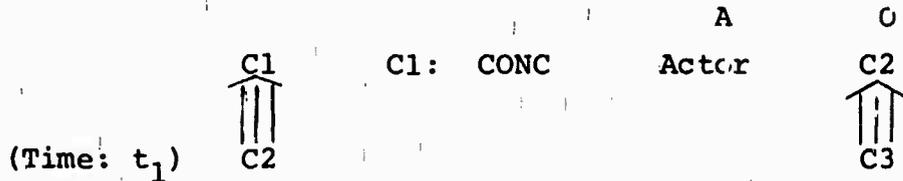
		A	O
		C1: CONC	
(TIME: t_1)	C2	C2: HIT	Fred John
(TIME: t_0)		C3: unstated	

That is, "Fred hit John" at Time t_1 but C3 had occurred at Time t_0 ($t_0 < t_1$).

Thus for "needlessly" we have in our dictionary:

Needless, needlessly:

The modified conceptualization is C2 in the following structure:



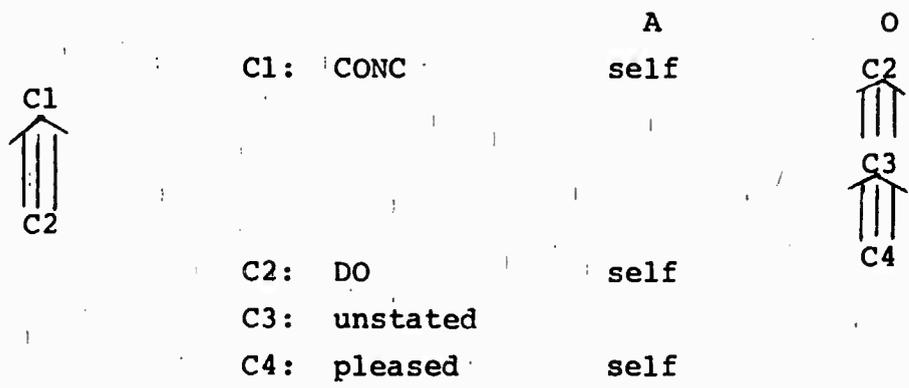
and in sense 1 add:

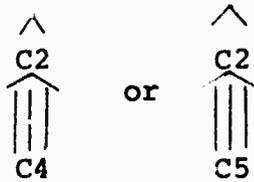


for sense 2 add:



The adverb "stupidly" refers to the same concepts that "needlessly" refers to. "Stupidly" in effect says, "I CONC-ed that goal X would result" and that either "but X didn't result" or "X did result but it did not make me happy". The first of these is of course exactly sense 1 of needlessly. the second sense refers to the principle that any action that one intends to do, one expects will please one in some way. Thus we actually have for "intend to do C2":

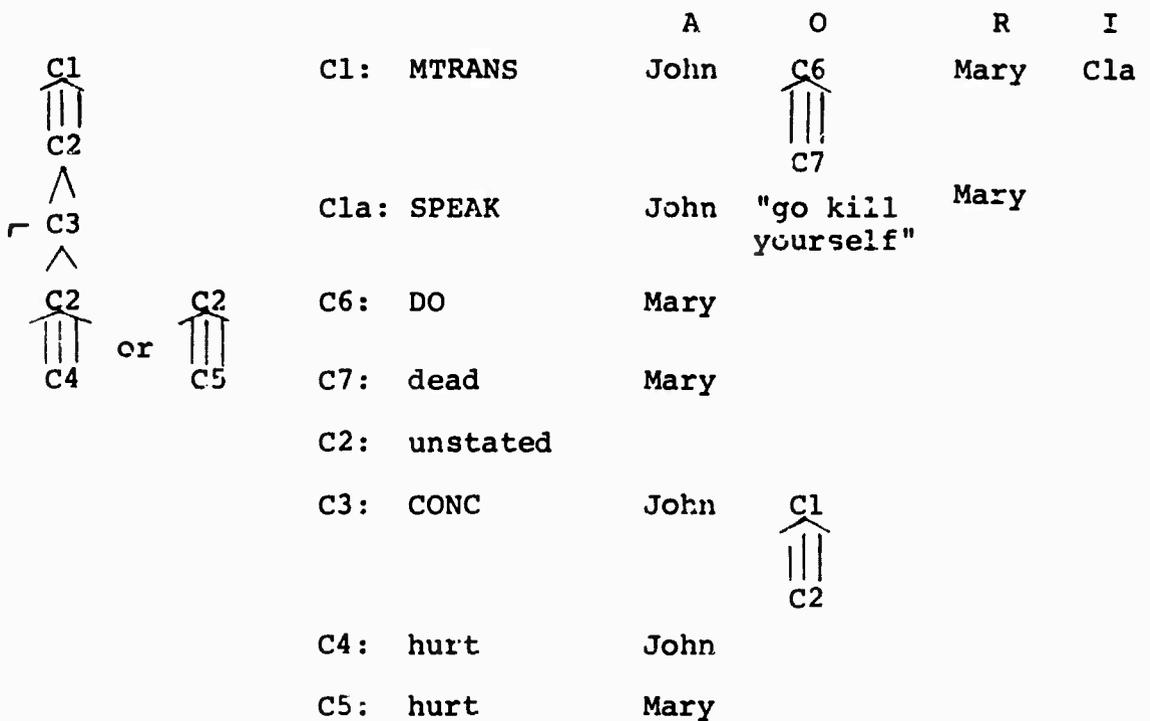




A
C4: hurt Actor

A
C5: hurt Object

Before we give the analysis of "thoughtlessly" in sentence (9), it is necessary to point out that "tell" is written conceptually as MTRANS (move an idea) by means of SPEAK words. Thus, our analysis for (9) is:



This complicated structure means "John communicated to Mary that she should kill herself caused C2, and John didn't think about the fact that it would cause C2, but C2 either

caused John to be hurt or Mary to be hurt or both". What bad thing may have actually occurred has not been stated. It might have been that John felt bad or Mary felt bad or that Mary killed herself.

It should be pointed out here that although the above structure is rather complex, it is in the nature of human language to be telegraphic. It is therefore to be expected that when the underlying conceptual structures that have been referenced in a shorthand manner are made explicit, that they will be complicated.

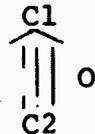
Another class of conceptually interesting adverbs are illustrated by the following sentences:

- (11) John punched Fred wrongfully.
- (12) John hit Fred unjustly.
- (13) Fred drove his car illegally.

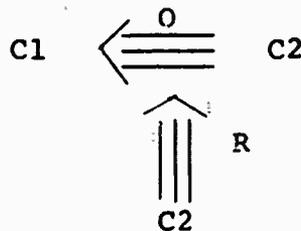
Whereas these sentences express a judgment on the part of the speaker, they are really more than just that. Actually, the judgment that is expressed is in terms of a belief. The belief that is referenced in (11) and (12) is a variation of the "revenge" belief. It has to do with a notion of justifiable cause. We have seen that the "revenge" belief can be used as a reason for an action. The adverbs "wrongfully" and "unjustly" refer to the fact that in the mind of the speaker (as opposed to the actor) the justification-reason that was responsible for the action on the part of

ceptualization under discussion. Another sense of "wrongfully" refers to the fact that such a justification belief does not exist in the mind of the speaker. This sense brings up the problem of "ought beliefs", which are different from the "justification-beliefs" that we have been discussing.

An "ought-belief" expresses a connection that has nothing whatever to do with the "vengeance" type belief but rather expresses a moral judgment of a sort (see Price [1969] for a discussion of this). We claim then that another reason for doing something is the existence of an ought belief. An "ought-belief" is of the form:



where C1 and C2 are conceptualizations and O indicates that when C1 occurs, C2 ought to occur. "Ought-beliefs" can, of course, serve as reasons for actions (particularly C2). So we often might have:



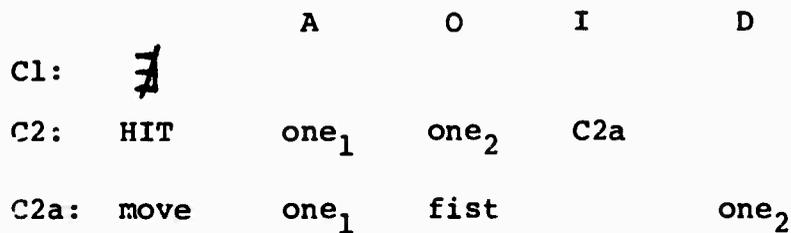
(That is, "since one ought to do C2 when C1 I did C2".) An example of an ought belief is expressed in sentence (14):

(14) John gave Mary flowers because he said he would.

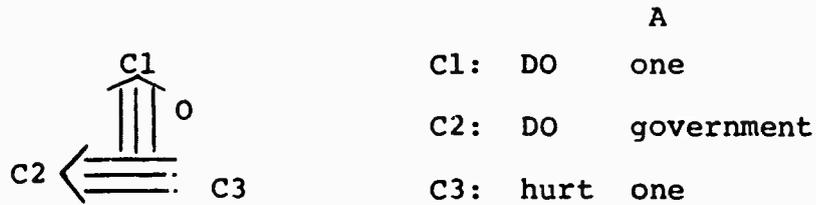
In (14) we have the "ought-belief" and action:



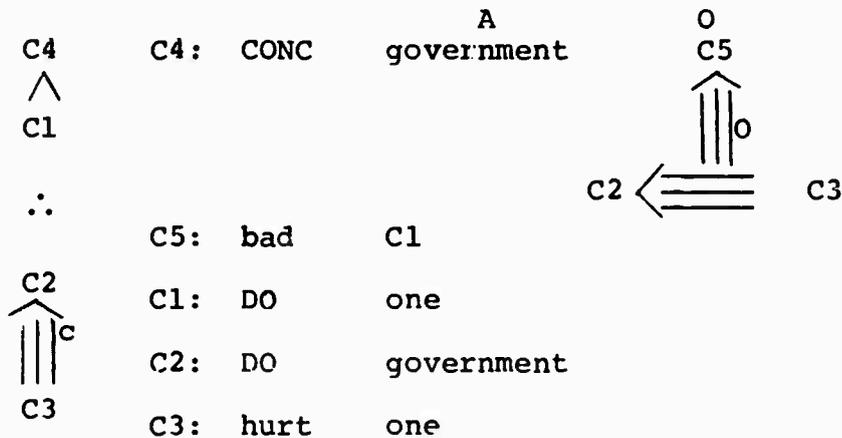
With this notion of an "ought-belief", we can explain another sense of "wrongfully", referring to an "ought-belief" which negates the actor's right to do a given action. So one sense of (11) is paraphrased by: Since nothing oughts punching when John punched Fred it was an instance of action without justification which is bad. This then reduces to a kind of logical paradigm:



(that is, there does not exist C1 such that it will ought C2).



Thus "illegally" is a statement that the referenced action fits the above paradigm as C1. That is, we have both the above paradigm and the one for "wrongfully". Thus we have:



This, then is "illegally" (i.e., the government thinks that a certain action is bad and is a reason for punishment. The referenced action (by illegally) is an instance of such a bad action and therefore the government could do something (c) to retaliate.) Note that the "vengeance" structure is present in "illegally" (as the object of CONC).

Another class of adverbs refer to the mental state of the subject. Often the use of these adverbs cause certain inferences to be made. One adverb whose inferences are within the vengeance paradigm is "angrily". Consider

sentence (15):

(15) John hit Fred angrily.

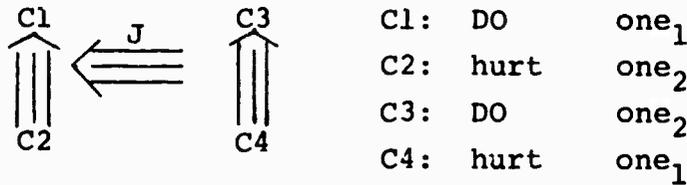
Is this any different than "vengefully"? Actually it is, in that it is vengefully plus anger. But the vengeance is, in this case, purely an inference and is not implicit within (15). That is, we have:

(Time: t_1) C1: HIT John Fred

(Time: t_2) C2: angry John

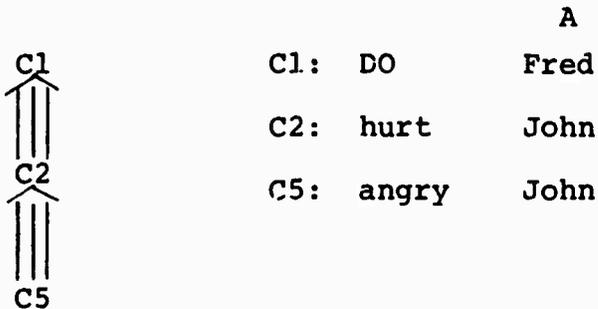
where $t_1 = t_2$

People make an inference here that since C1 causes the object to be harmed and since one is angry when one feels oneself to have been harmed, then C1 above is equal to C3 of the vengeance belief and C2 is a consequent of the C2 of the vengeance belief. We thus have the vengeance belief:



and have identified John as one₂ and Fred as one₁ of C4. The inference is then natural that Fred must also fit as actor in C1; that is, that he must have done something that hurt John that got him angry.

What we are claiming here is that since the vengeance belief paradigm is so common that when people see conceptualizations that fit into pieces of it, they often infer (perhaps incorrectly) that the other pieces are present also. We thus claim that a natural inference from sentence (15) is:



4. Conclusion - The point of this paper has been to show that many adverbs in English are beliefs conceptually. The common idea that adverbs "modify" verbs can only be transferred to the conceptual level if the verb that is in use is an ACT conceptually and if the modifying adverb refers to a particular aspect of that ACT. Thus, if we have "run fast" or "hit hard", we have a case of an adverb syntactically that modifies an action conceptually. But, many adverbs do not conceptually modify an ACT but rather refer to some beliefs that are present in the memory structure of the speaker which in effect modify the entire conceptualization underlying the given sentence.

The attempt here has not been to give all possible

analyses of each adverb presented. Certainly, each example adverb given has other meanings with which we have not dealt here. The example sentences often can be interpreted in an alternative fashion from the one chosen. The point here is to show how such adverbs should be dealt with rather than to exhaustively deal with each of them.

It is also true that we have diverged considerably from the usual manner in which linguists deal with sentences. An understanding theory such as we have been interested in is not concerned with the grammaticality or ungrammaticality of sentences. We are simply interested in assigning conceptual structures to input sentences. We claim that it is the responsibility of linguists to provide formal rules for assigning such conceptual structures as a proper domain of linguistic theories.

We claim here that it is the proper domain of linguistics to explore this important relationship between language and beliefs, and that such an explanation is more easily done within a "performance" theory of language understanding than within the traditional competence framework.

REFERENCES

1. Abelson, R. and Carroll, J. 1965 "Computer Simulation of Individual Belief Systems". American Behavioral Scientist vol. 8.
2. Chomsky, N. 1965 Aspects of the Theory of Syntax. M.I.T. Press.
3. Colby, K., Weber, S. and Hilf, F. 1971 "Artificial Paranoia" Artificial Intelligence vol. 2.
4. Lakoff, G. 1970a "Adverbs and Modal Operators". Phonetics Lab, University of Michigan (mimeographed).
5. Lakoff, G. 1970b "Adverbs and Opacity". Phonetics Lab, University of Michigan (mimeographed).
6. Price, H. 1969 Belief George Allen and Unwin Ltd.
7. Schank, R. (in press) "Conceptual Dependency: A Theory of Natural Language Understanding". Cognitive Psychology.
8. Schank, R., Goldman, N., Rieger, C. and Riesbeck, C. 1972. "Primitive Concepts Underlying Verbs of Thought". Stanford Artificial Intelligence Memo 162, Computer Science Department, Stanford, California.