Title
Identification of memory processes underlying text understanding : a research proposal

Authors
Granger.Jr., Richard H.
Meehan, James R.

Publication Date
1981

Peer reviewed

IDENTIFICATION OF MEMORY PROCESSES UNDERLYING TEXT UNDERSTANDING:

A Research Proposal

Richard H. Granger, Jr.
and
James R. Meehan
Artificial Intelligence Project

Technical Report #163

January 1981

This work was supported in part by the Naval Ocean Systems Center under contract N6601-80-C-0377.

1.0 Introduction

The UCI Artificial Intelligence Project has as its goal the creation of computer programs that model human thought processes. Our primary interest is in the cognitive mechanisms people have for using language. We are interested in building a theory of human natural language processing and in testing that theory by constructing computer models. Thus the goal of our project is twofold: we want to develop a theory of human cognition, and to produce working programs capable of performing useful tasks that involve natural language.

In much early work on language processing, syntax was considered a central aspect of the nature of language, and syntactic parsers became very important. There is a large literature on parsing techniques for a variety of formal languages (context-free, LR(k), etc.). The mechanisms were very well understood, and efficiency became a matter of prime concern. The bulk of recent work in natural language processing in the fields of Artificial Intelligence and Cognitive Psychology [36,34,35,44,43,45,25,19,18,20,10,7,4,33,28,39] has regarded understanding to be a mapping of natural language text into a language-free representation of the meaning of that text. Hence, much of the recent work in natural language has focused on the representation of meaning, the structure of knowledge, the organization of that knowledge in memory, and the process of producing these representations from natural language input.

Many current process models of language understanding [3,6,7,12,15,20,30,31,32,34,35,37,42,44,45] are based primarily on a process called "expectation-based understanding" [31]. Expectation-based understanding involves the reader's generating alternative expectations about the possible meanings of a text, eventually producing a single, unambiguous meaning. Generating expectations from a piece of text helps in discriminating among the possible readings of subsequent text: readings which correspond closely to an existing expectation are preferred over those not predicted. Reciprocally, subsequent input helps to select a single unambiguous representation from the initial large array of alternative expectations.

The process of expectation-based understanding can account for human readers' ability to arrive at a single meaning representation for a sentence in context which would be highly ambiguous in isolation. For example, consider the following:

[1] Kathy and Chris were playing golf. Kathy wanted to let Chris win the game. She hit a shot deep into the rough.

We infer from the last sentence that "she" is Kathy, and that she hit her ball into the rough on purpose, since that will further her stated goal of letting Chris win the game. The reason we can make these inferences is that the previous sentences have suggested expectations, one of which can be stated as follows:

> Kathy will perform some action in service of her goal of letting Chris win the golf game.

When she then hits her ball into the rough, that action is checked against this expectation (among others), and it is seen that the action can be interpreted as being in service of Kathy's goal, and hence it is so interpreted. Without such expectations, the statement "She hit a shot deep into the rough" would more likely be interpreted as being an accidental shot, rather than being intentionally performed.

Hence, understanding based on expectations (and subsequent substantiation of those expectations) is helpful in accounting for how people arrive at an interpretation of a story. If all stories could be understood by the reciprocal processes of expectation and substantiation, then future work in natural language processing could focus exclusively on issues of the representation of knowledge and the organization of that knowledge in memory. Unfortunately, this turns out not to be the case. First of all, people are able to understand wholly unexpected events in a story, which a strictly expectation-based understander would fail to interpret correctly. This is especially clear in the case of a story which contains a surprise ending, as in the following version of one of Aesop's fables:

> [2] The crow was sitting in her tree holding a piece of cheese, when along came the fox. "Sing for me, Crow," implored the fox, "for you have the most lovely voice in all the world." The crow was much impressed by this flattery, and she opened her mouth to sing. When she did so, the cheese fell out of her mouth, and the fox ran away with it, laughing as he ran. The cheese was what he wanted all along.

When we first read the fox's request that the crow sing, we can only infer that the request is genuine, and that he wants to hear her singing. We may suspect that the fox has some hidden plan in mind, but we certainly cannot tell what that plan is from the story at that point. When we eventually see that the cheese falls and the fox grabs it, we can then recognize that the fox's actual goal was to get the crow to drop the cheese. Had this goal been stated at the beginning of the story, we would have made different inferences from the fox's actions.

The experience of conscious surprise while understanding a story indicates that the reader has generated some predictive inference or expectation, and has subsequently decided to invalidate that expectation and supplant it with a different inference that was not previously predicted. For example, consider the following version of the old Henny Youngman joke:

> [3] I took my wife to the Bahamas, but unfortunately she found her way back home.

When we hear the first part of this sentence we infer that the speaker took his wife with him on a vacation. However, the second half of the example says that his wife found her way back home, and implies that he didn't want her to do so. Upon hearing this, we abandon our initial inference, and infer instead that the speaker took his wife to the Bahamas so that she would stay there and be out of his proximity from then on.

If the example instead said "I took my wife to the Bahamas, and we had a wonderful vacation", the natural inference predicted from the first part of the example is substantiated by the subsequent part of the statement, and no surprise occurs.

Even when a story does not cause us to experience conscious surprise, a reader may have to make an inference and then change it in light of subsequent input. This can occur when a non-standard activity is performed in an otherwise stereotypical setting. For instance, consider the following altered version of the above "golfing" example:

> [4] Arnie and Jack were playing golf. Arnie hit a shot deep into the rough. He wanted to let his good friend Jack win the game.

(Note that the second and third sentences have been transposed from the previous version.) In this example, the most natural default inference from the first sentence is that both Arnie and Jack had the goal of winning the game, or at least playing well. In light of this inference, Arnie's action of hitting into the rough is most likely to be an accidental action which hinders his goal. Finally, the third sentence explicitly contradicts our initial default guess of Arnie's goal, and concomitantly requires us to re-evaluate our inferences about his action: hitting a shot into the rough can now be assumed to have been <u>intentionally</u> performed by Arnie, in service of his stated goal of wanting to let Jack beat him.

Expectation and substantiation are but two of the processes that may underlie the ability to understand natural language. It is the long term goal of our research project to determine what other types of basic processes there may be. The purpose of this proposal is to address the problem of basic knowledge processes on three fronts:

1. Identifying what basic knowledge processes underlie cognitive tasks such as natural language understanding in addition to those we have already hypothesized.

2. Developing the theory of the basic knowledge processes of which we are currently aware to the point where we can understand their uses and limitations.

3. Incorporating this developing theory into computer programs with the intention of both evaluating the theory and extending the capabilities of our existing language understanding programs.

We feel that we have already made considerable but unexhaustive progress towards developing a theory of basic knowledge processes. We have identified a preliminary set of five processes beyond the expectation-substantiation mechanism. The bulk of our effort will be to investigate the nature of each of these processes and their interrelationships and dependencies as they are applied during the understanding process. In order to effect this investigation, we plan to build computer models of these understanding mechanisms, with an eye to observing their separate contributions and their interactions in the operation of an integrated understanding system. Furthermore, we hope to investigate how these knowledge processes might be applied in domains other than language understanding, especially problem-solving tasks such as game playing (e.g., chess end-games) and medical diagnosis (such as in a MYCIN-like system).

We have seen from the examples above that expectation and substantiation are insufficient to account for the processing of certain stories. In the following sections we will briefly introduce the five additional knowlege processes that we intend to investigate. They are described in greater detail in section 2.

1.1 Correcting erroneous inferences

When we write or speak, we omit great quantities of information that we expect the reader or hearer to infer. For example: "Little Johnny got an F on his report card. He was afraid to show his father." There is a long chain of inferences that logically connect those two sentences, and there are many models of language that decribe how to fill in such gaps. But there is rarely any guarantee that the inferences are correct, merely that they are plausible.

Previous theories of text understanding [7,34,35,12,42] have assumed that people make inferences while reading. Furthermore, all of these theories acknowledge that an inference made while reading might turn out to be wrong. However, no previous theory of text understanding has included a method of either recognizing when a previous inference is wrong nor a method for replacing one inference by a subsequent one based on new information in a story. The knowledge process of <u>supplanting</u> an initial inference with a new one can account for this ability of understanders to "change their minds" about the interpretation of a story.

1.2 Resolving loose ends

Rarely is a piece of text so constructed that every new idea is developed and completed before any other new ideas are mentioned. In short texts, the expectations that are pending are usually resolved quickly, but in longer texts, the problem of pending expectations, or "loose ends," can be severe. When we finish reading a text, or some section of it, we notice that some of our expectations were not completed or not fulfilled.

Perhaps the simplest case where expectations cannot be immediately resolved is a narrative describing the actions of several characters, in which the presentation of events is chronological. After hearing about one character for a while, we hear about some of the others, but we expect that the story will get back to the first character at some point. If not, we notice the lack of completeness.

A more complicated case would one where the expectations were structural, not sequential. If a murder mystery were to end before the murderer was identified, we would certainly notice.

Of course, not all details <u>require</u> completion, and part of our investigation of loose ends will include a distinction between "hot" loose ends which require resolution, and "cold" loose ends which may be left dangling even after the story is complete.

1.3 Assimilating new implication rules

Through the acquisition of new information, the inferences that the reader draws may change. A new fact in a story may cause the reader to generate inferences from certain story events that would not otherwise have been made. For example, we may have trouble understanding the connection between the two events in the following story:

> John was taken to Hoag Hospital with a case of hepatitis. The nurses fed him some Jameson's Irish Whisky.

unless the story then tells us that

> It had recently been discovered that Jameson's Irish Whisky can cure hepatitis.

This last statement is a fact which causes us to add to our memories a new "implication rule," which in turn gives rise to a previously unknown inference from certain events such as drinking Jameson's.

1.4 Understanding anomalous situations

When we fail to understand the connection among the statements in a story, and no new implication rule is presented which allows us to generate any bridging inferences, then the story situation can be described as an anomalous situation. In such a case, a reader may attempt to generate a tentative new implication rule that will give rise to an inference that would allow the story statements to be connected. For example, in the previous story about Hoag Hospital, in the absence of any information about the healing powers of Jameson's whisky, we might instead have assumed that the nurses were simply trying to comfort a dying man, in spite of the fact that alcohol might be detrimental to the patient's condition. We call the process of generating such an implication rule "forcing an interpretation."

1.5 Recognizing contextual significance

Sometimes a reader will choose to simply ignore a given piece of text, either because it presents him with an anomaly he chooses not to solve, or because the passage seems irrelevant to the rest of the text. The ability to distinguish between central and peripheral parts of a text is evidenced by people's tendency to consistently remember certain parts of a text over others, and their ability to summarize a text, leaving out the details. These abilities can be attributed to a process of constructing a story representation in such a way as to assign relevance to certain parts of the story while relegating other parts to a lower position in the hierarchy of events.

2.0 Details of the proposed research

We have briefly specified five knowledge processes beyond expectation and substantiation. It is our intention to investigate each of these five processes in depth, and to study their interactions within a complete understanding system. We have previously developed a computer program called ARTHUR [16,17], which incorporates our early theories of some of these processes, especially that of correcting erroneous inferences. The rest of this section will describe each of these processes in more detail, including some explanation of ARTHUR's operation. It is our intention to build a computer understanding system whose capabilities exceed those of ARTHUR. Such a program requires extending the theory of basic knowledge processes that we currently possess.

2.1 Correcting erroneous inferences

A normal, necessary part of the process of understanding a piece of text is to make plausible inferences along the way as to the intentions of the characters in the text. These intentions can often be predicted once the context has been established. But for any given context of intentionality, there are many possible explanations, and human readers are good at picking the most accurate one at the outset.

There already exist computer models [12,42] that are capable of understanding such texts by inferring the correct context (defined in terms of Schank and Abelson's scripts, plans, and goals [35]) and matching further input with predictions from that context.

But it is not always possible to make the correct choice at the beginning, and readers have the ability to revise their understanding by correcting their erroneous inferences as they read the text. We have previously studied several kinds of the erroneous inferences that can be corrected during understanding, and we have embodied the theory into a computer model called ARTHUR (A Reader THat Understands Reflectively) [16,17].

2.1.1 Goal errors

When we read about a certain action, we can infer the goal that motivated that action. In general, we prefer the most parsimonious explanation; that is, given two actions, we would prefer to find one goal that explains both actions rather than two goals, one explaining each event. After reading an initial sentence, we may infer a motivating goal that subsequently is less preferable on grounds of parsimony, so we revise our analysis of the text to include the new goal.

For example, consider this simple story:

Mary picked up a magazine. She swatted a fly.

The first sentence causes ARTHUR to generate the plausible inference that Mary's goal is to read the magazine for entertainment, since that is stored in ARTHUR's memory as the default use for a magazine. ARTHUR's internal representation of this situation consists of three items: a goal (being entertained), an event (picking up the magazine), and an inferential path connecting the event and goal (reading the magazine). (The upper-case text below is actual output from the ARTHUR program.)

```
:CURRENT EXPLANATION-TRIPLE:
GOAL:    (ENJOY-ENTERTAINMENT (PLANNER MARY) (OBJECT MAGAZINE))

  EVENT:   (GRASP (ACTOR MARY) (OBJECT MAGAZINE))
   PATH:    (READ (PLANNER MARY) (ACTOR MARY) (OBJECT MAGAZINE))

:NEXT SENTENCE:
She swatted a fly.

:CD:
(PROPEL (ACTOR MARY) (OBJECT NIL) (TO FLY) (FROM NIL))
```

ARTHUR has now read the second sentence and displays the Conceptual Dependency representation of Mary's action: she struck a fly with an unknown object. It next infers Mary's goal is hitting the fly.

```
:BOTTOM-UP INFERENCE PATH:
(CHANGE-PHYS-STATE (PLANNER MARY) (ACTOR MARY) (OBJECT FLY)
                   (DIRECTION NEG) (INSTRUMENT NIL))
```

ARTHUR infers that the most plausible explanation for Mary's action is that she intends to damage the fly (CHANGE-PHYS-STATE (DIRECTION NEG)). It now tries to connect that goal with the previous goal of being entertained.

:ATTEMPTING TO CONNECT TO EXISTING GOAL CONTEXT
(ENJOY-ENTERTAINMENT (PLANNER MARY) (OBJECT MAGAZINE))

:CONNECTION FAILS

ARTHUR's initial goal inference (that Mary plans to entertain herself by reading the magazine) has failed to explain Mary's subsequent action of swatting a fly. ARTHUR has now recognized that the new event cannot be adequately explained in terms of the initial goal inference. It therefore must generate an alternative goal inference on the basis of the new event, and ensure that this new goal inference can also account for the previous story events.

:GENERATING PLAUSIBLE GOAL INFERENCE:
(PHYS-STATE (PLANNER MARY) (OBJECT FLY) (VAL -10))

ARTHUR now generates an alternative goal inference on the basis of Mary's new action: she may plan to destroy the fly (PHYS-STATE (VAL -10)).

:ATTEMPTING TO SUPPLANT PREVIOUS PLAUSIBLE GOAL INFERENCE

ARTHUR infers that this new inference may also serve to explain Mary's previous action of getting a magazine, in place of the initial "entertainment" inference which currently is the explanation for that action.

:CONNECTION SUCCEEDS

The new inference can explain her previous action, ...

:SUPPLANTING INITIAL INFERENCE PATH:
(DELTA-CONTROL (PLANNER MARY) (ACTOR MARY) (OBJECT MAGAZINE)

... via an inference that she picked up the magazine as a precondition (DELTA-CONTROL) to using it for some other plan (which in this case was changing the physical state of the fly).

```
:AUGMENTING INFERENCE PATH:
(CHANGE-PHYS-STATE (PLANNER MARY) (ACTOR MARY) (OBJECT FLY)
                   (DIRECTION NEG) (INSTRUMENT MAGAZINE))
```

As a side-effect, ARTHUR now notices that the magazine must have been the INSTRUMENT in Mary's plan to damage the fly.

```
:UPDATING EXPLANATION-GRAPH
```

ARTHUR now has a representation of the two events in the story in terms of a single goal (destroying a fly).

```
:CURRENT EXPLANATION-TRIPLE:
GOAL:    (PHYS-STATE (PLANNER MARY) (OBJECT FLY) (VAL -10))

  EVENT:   (GRASP (ACTOR MARY) (OBJECT MAGAZINE))
   PATH:    (DELTA-CONTROL (PLANNER MARY) (ACTOR MARY)
                           (OBJECT MAGAZINE))

  EVENT:   (PROPEL (ACTOR MARY) (OBJECT MAGAZINE) (TO FLY))
   PATH:    (CHANGE-PHYS-STATE (PLANNER MARY) (ACTOR MARY)
                               (OBJECT FLY) (DIRECTION NEG)
                               (INSTRUMENT MAGAZINE))
```

This representation says that Mary wanted to destroy a fly (PHYS-STATE), so she planned to physically damage the fly (CHANGE-PHYS-STATE), which required her first to get control of some object (DELTA-CONTROL) that could be used to do so. These two plans were carried out by getting hold of a magazine (GRASP) and then hitting the fly with it (PROPEL).

```
:SUPPLANTED INFERENCES:
GOAL:    (ENJOY-ENTERTAINMENT (PLANNER MARY) (OBJECT MAGAZINE))
  EVENT:   (GRASP (ACTOR MARY) (OBJECT MAGAZINE))
   PATH:    (READ (PLANNER MARY) (ACTOR MARY) (OBJECT MAGAZINE))
```

The initial goal inference, that Mary may have wanted to read the magazine for entertainment, is no longer considered by ARTHUR to be an explanation for having picked up the magazine.

:UNDERSTADING PHASE COMPLETED
:READY FOR QUESTIONS

>Why did Mary pick up a magazine?

AT FIRST I THOUGHT IT WAS BECAUSE SHE WANTED TO READ IT, BUT ACTUALLY IT'S BECAUSE SHE WANTED TO USE IT TO GET RID OF A FLY.

> The question asks for ARTHUR's inferred goal inference underlying Mary's action of GRASPing a magazine. The answer is generated on the basis of ARTHUR's supplanted inference about this action (READ) and the currently active inference about the action (CHANGE-PHYS-STATE). (The English generation mechanism used is outlined in Granger's thesis [16].)

2.1.2 Plans and Scripts

A second type of error is one where we infer the correct goal but the incorrect plan for achieving that goal. Consider this story:

> Carl was bored. He picked up the newspaper. He reached under it to get the tennis racket that the newspaper had been covering.

In this example, ARTHUR correctly infers the goal of the story character, but erroneously infers the plan that he is going to perform in service of his goal. By knowing that the <u>function</u> of a newspaper is to be read, ARTHUR first infers that Carl will read the newspaper to alleviate his boredom, but this inference fails to explain why Carl then gets his tennis racket, so ARTHUR attempts to supplant the initial goal inference. However, in this case ARTHUR knows that that goal was <u>correctly</u> inferred, because it was implicitly stated in the first sentence of the story (that Carl was bored). Hence ARTHUR decides instead that it erroneously inferred the plan for satisfying that goal, namely reading the newspaper. By knowing that the function of a tennis racket is to play tennis, ARTHUR infers that Carl planned to alleviate his boredom by playing tennis. Now we have to explain how picking up the newspaper fits in with playing tennis. Instead of using the newspaper as a <u>functional</u> object (in this case, reading material), ARTHUR considers whether the newspaper might be an <u>instrumental</u> object for the <u>precondition</u> of getting to the tennis racket (which is itself a precondition to playing tennis), and indeed ARTHUR figures out that the newspaper must be moved in order to get the racket. Thus, correcting this erroneous plan inference required ARTHUR to re-evaluate its inference about the intended use of a functional object.

2.1.3 Causal state changes

A third kind of error is that in which the goal is correct, the plan is correct, but a subsequent action actually hinders the plan. Consider this example from section 1:

> Kathy and Chris were playing golf. Kathy hit a shot deep into the rough.

Since Kathy is playing golf, ARTHUR infers that she plans to win or at least to play well, but it understands that hitting a ball into the rough hinders either of those goals. (ARTHUR can decide whether an action achieves, helps, hinders, or thwarts a goal.) ARTHUR infers that the action is accidental; that is, it is an action which causally results in some state which may have a negative effect on her goal.

This situation differs from the two previous stories in that ARTHUR does not change its mind about its inference of Kathy's goal or her plan for achieving that goal. Subsequent input might force it to, however. If the next sentence were:

> Kathy wanted to let her good friend Chris win the game.

then ARTHUR would have to recognize that the initial goal inference was erroneous, and to supplant it by the inference that Kathy actually intended to lose the game, not win it.

2.1.4 Proposed extensions

While the process of correcting erroneous inferences is by far the best developed of the five processes we propose to study, more work remains to be done.

Our first effort will be to sharpen our definition of what "accidental" means, by categorizing the events that directly or indirectly contradict a goal-predicted event. It is not sufficient merely to notice that an event does not match a prediction. Non-matching events might be totally unrelated to the current goals. On the other hand, the contradiction might be highly indirect. For example, consider the following sentences: "John wanted to fly to New York. He lost his VISA card." Contrast that with: "John wanted to fly to New York. He used his VISA card." Although the examples are lexically nearly identical, we understand that losing the card is an accident, hindering the goal of getting to New York, whereas using the card is an intentional action in service of the goal.

2.2 Loose ends

When the correct context inference for a story remains unknown until some significant fraction of the story has been read, the story can be thought of as a "garden path" story. This term is borrowed from so-called garden path sentences, in which the correct representation of the sentence is not resolved until relatively late in the sentence. We will call a garden path story any story which causes the reader to generate an initial inference which turns out to be erroneous on the basis of subsequent story events. Obvious examples of garden path stories are those in which we experience a surprise ending, e.g., mystery stories, jokes, fables. Since ARTHUR operates by generating tentative initial inferences and then re-evaluating those inferences in light of subsequent information, ARTHUR understands simple garden path stories.

Not all garden path stories cause us to experience surprise. For instance, recall the following simple example:

Mary picked up a magazine. She swatted a fly.

Many readers of this story do not notice that Mary might have been planning to read the magazine, unless that intermediate inference is pointed out to them. Hence we hypothesize that the processes involved in understanding stories with surprise endings must differ from the processes of understanding other garden path stories. Therefore a shortcoming of ARTHUR's understanding mechanism from the standpoint of psychological plausibility is that ARTHUR does not differentiate between stories with surprise endings and other garden path stories.

We are currently developing a more sophisticated theory of the processes underlying story comprehension which accounts for the differences between the processing of garden path stories with surprise endings and those without. We intend to incorporate this new theory into a new version of ARTHUR (called "Macro-ARTHUR"). MacARTHUR will differentiate between "strong" default inferences and "weak" tentative inferences when generating an initial context inference. When a strong initial inference is generated, then MacARTHUR will consciously "notice" this inference being supplanted and consequently experience surprise that the inference was incorrect. Conversely, when the initial inference is weak, MacARTHUR will not commit itself to that inference, but rather will choose to keep around other possible alternatives. We call this latter situation leaving a "loose end" in the story representation being constructed. Once MacARTHUR has left a loose end in its current story representation, then resolving that loose end will involve only further specification of the initial tentative set of inferences, rather than supplanting a single strong inference.

We hypothesize that a human understander will experience conscious surprise accompanying the process of supplanting a previous inference, while the process of narrowing down a set of tentative inferences (resolving a loose end) will not result in conscious surprise. Few psychological studies of the cognitive phenomena underlying human story understanding have so far provided few data relevant to the question of when readers process consciously versus unconsciously. (Studies include Thorndyke [39] and Norman and Bobrow [46].) Our new understanding mechanism will have available to it both the process of correcting erroneous inferences and the process of leaving and subsequently resolving loose ends while reading text. These two processes seem to bear on the issue of conscious vs. unconscious processing, so it is our hope that study of the operation of our understanding model may shed some light on this currently open question in cognitive science.

2.3 Assimilating new implication rules

Stories sometimes contain wholly novel information, which must be understood and assimilated in order to enable understanding of the rest of the story. For example, consider the following:

[5] John was hungry. He went outside and collected some ants. John was the "human anteater" at the circus.

Q1) Why did John get some ants?
A1) He wanted to eat them.

Story [5] contains an explicit statement saying that John regularly eats ants. Generating answer A1 requires the reader to use this statement about John's eating habits to infer a novel connection between his action of getting ants and his goal of satisfying his hunger.

This story leads the reader to assimilate a previously unknown fact about the world -- that a particular actor may eat ants to satisfy his hunger. This new fact is not simply a static piece of information. Having acquired this new knowledge, the reader will now draw different inferences from certain events than he previously would have generated.

For example, once the reader knows this fact about John and ants, then whenever the story next says that John got some ants, the reader can infer that John may plan to eat the ants. This inference is not one that would have been made prior to the introduction of this new fact. Hence, this fact can be thought of as a new _implication rule_ for generating inferences during understanding. Understanding a new implication rule involves assimilating it into memory such that it can be used to generate previously unavailable inferences during understanding.

Incorporating this ability into our model of understanding involves a representation issue as well as a processing problem. We normally think of inferences as being expressed in the form "If A then B"; e.g., "If John is hungry then he may eat ants." However, many inferences are expressed as state-information. When we say "mushrooms are poisonous," for example, what we mean is that if you eat one, you will probably become ill. So what looked like a state (MUSHROOMS +POISONOUS) is actually better represented as a rule (IF (INGEST MUSHROOMS) THEN -HEALTH). In fact, it may be that most states are equivalent to a set of inferences via implication rules in some useful and non-trivial way.

2.4 Understanding anomalous situations

A situation may arise in a story that appears anomalous to the reader; that is, the reader is unable to make sense of the events that occur in the story. Consider just the first two sentences of example [5] above:

> [6] John was hungry. He went outside and collected some ants.

In the previous version of this example, it was explicitly stated that John ate ants for a living in the circus. Understanding that story required using that fact to generate a new inference connecting John's hunger to his getting ants. In the current example, however, no such explicit new implication rule is given. Yet readers of this story typically answer, when queried, that John probably planned to eat the ants. Giving this response implies that the readers have somehow generated and then applied the very implication rule that was explicitly supplied by the previous example. People typically attempt to connect up the events in a story by finding some unifying thread that serves to tie together otherwise disconnected events. In order to do so in example [6], the reader creates an implication rule which allows a bridging inference to be generated for this example. We call the process of generating such an implication rule "forcing an interpretation."

2.5 Recognizing contextual significance

Not all events are created equal. It is important to be able to distinguish between details of a text that are crucial to understanding and those that are not. In effect, this amounts to knowing the <u>bounds of relevance</u> for any event. While we customarily think of events in a text as being goal-driven and somehow all connected, we need to know the context of each event, that is, the goal that motivates each event, so that we understand it relevance to the text in general.

Furthermore, there are hierarchies of goals -- only some goals are crucial to the action in the text. In general, high-level goals are distinguished from low-level goals by the effect their achievement (or lack of achievement) has on the rest of the text, and an event inherits the significance of the goal that motivates it.

Schank and Abelson [35] describe a problem-solving hierarchy of goal types for social and physical acts. This is useful for measuring one type of contextual significance, namely, that of goals and subgoals. For example, an event within the context of a script is not more significant than the script is itself.

Another measure of significance comes from the type of text. The relevance of a particular event appearing in a romantic novel ("John walked into the room") may be quite different from the relevance of the same event in a murder mystery, and a good theory of significance must account for the different modes of processing such information.

One application of such a theory is the ability to summarize a text by representing only the significant goals and events. The detail included in the summary would be determined by a significance threshold in the understander. By varying this threshold, one might be able to produce the effect of skimming a text.

3.0 Applications outside text processing

While all the examples we have presented so far have been related to text understanding, we believe that these same five processes, along with the traditional expectation-substantiation mechanism, can be applied to other domains as well. In particular, they have relevance to expert systems and problem solving. To illustrate this, we will consider some issues raised by the application of the basic knowledge processes introdued in this proposal to an expert medical-diagnosis system, a chess player, and a mathematical theorem prover.

Medical diagnosis

1. Correcting erroneous inferences. Given the incomplete state of our knowledge about medicine, let alone what we are able to encode in a system, there will always be some uncertainty in the deductions we make. In fact, MYCIN, the best-known medical diagnosis system, incorporates certainty factors in all its calculations. Once an error is detected, the system must have the ability to trace the extent of the error and take appropriate corrective action.

2. Resolving loose ends. Does the final diagnosis account for all the symptoms? If not, should the diagnosis be abandoned, and will some other diagnosis, compatible with the first, be found?

3. Assimilating new implication rules. The likelihood of certain illnesses or a patient's reactions to therapy may change as information is added to the system. Assuming that the system will never contain an exhaustive set of diagnosis rules, it will be valuable to be able to add new such rules to the system in the form of "facts" about the world; e.g., "Patients with disease X often display symptom Y," which can be translated into an implication rule of the form "If the patient displays symptom Y, then consider the possibility of disease X."

4. Understanding anomalous situations. What should the system do with information that fits into no known pattern? Is there a reasonable guess available, or should the evidence itself be questioned?

5. Recognizing contextual significance. Certain findings are always important, but others depend on the particular illness at hand. For example, a particular symptom may be the deciding factor in diagnosing a certain disease, while it will be irrelevant to the diagnosis of some other diseases.

A game of chess

1. Correcting erroneous inferences. A player is often uncertain of the opponent's strategy, but he must be able to cope with wrong guesses and salvage the game if possible. For example, a player may have inferred that his opponent is pursuing a certain line of play, and react accordingly, only to discover later that he has made the wrong inference about his opponent's true aim.

2. Resolving loose ends. In planning future moves, a player must keep track of many details -- relative advantage, pieces under attack, and so on. In trying to understand the opponent's plans, does an explanation adequately account for all the moves so far?

3. Assimilating new implication rules. The rules for what to do in a particular situation may change depending on the state of the game, the player's information about the opponent's likely strategies, or the player's decision to change strategies in mid-game. It is unlikely, however, that a new implication rule will be specifically introduced to a player in the midst of a game, except in the circumstance of someone offering advice to the player, such as "If he maneuvers his rook onto an empty file, he may be planning to get behind your pawn structure."

4. Understanding anomalous situations. An opponent may make an unexpected move, or a player may find himself in an unforeseen situation. If no known implication rule suffices to account for the anomaly, the player may be forced to guess at a plausible explanation.

5. Recognizing contextual significance. Some moves may be made simply to mark time, but any move has the potential of being crucial, e.g., part of a sequence that ends with checkmate.

Proving a theorem

1. Correcting erroneous inferences. In proving a theorem, one may believe in implications that, when examined closely, turn out to be false. In following the proof of a theorem, a reader may make an incorrect guess as to the purpose of a particular statement or line of attack.

2. Resolving loose ends. Theorems are perhaps the most elegant case of having no loose ends. By definition, they are supposed to cover all the relevant cases and only those. A loose end in a proof may be fatal. However, not all of the implications of a particular

statement of a proof need be followed out to completion. That is to say that a proof may have other consequences in addition to the particular theorem it proved; it may, as a side-effect, contribute to the proof of some other theorem as well.

3. Assimilating new implication rules. Sub-proofs (lemmas) represent acquired information that may alter the proof strategy. If a lemma is introduced in the course of a proof, it is expected that the lemma will be used immediately.

4. Understanding anomalous situations. Anomalies can range from surprises to contradictions. A reader (e.g., student) who fails to understand a particular part of a proof may simply take it on faith that the proof is correct; a reviewer of a journal article containing a proof, on the other hand, may reject a proof that does not make explicit any information required to bridge a gap between consecutive statements of the proof.

5. Recognizing contextual significance. Mathematical proofs are so highly structured that the significance of any particular statement is assumed to be known; that is, its relevance to the entire process is often explicit. However, a summary can often be made of a proof in which the key steps are made explicit, under the assumption that any intermediate steps can be inferred independently by the reader.

4.0 Research method and facilities

Language research in Artificial Intelligence begins in the same way as it does in other, older disciplines -- choosing a problem, focusing on certain aspects of the problem, and finally asking questions that are precise enough to be answerable within the resource limits of the research project. At that point, AI diverges from other paradigms by restricting the possible answers to computably-precise process models. That is, to any theory describing what information people may have, we must attach an explanation for the ways in which that information is manipulated in performing the language task in question. The resulting theory takes the form of a computer program that solves the problem.

Building these program-models is a long, interactive and cyclic process of specifying the theory, expressing it in a programming language, and testing. With the current technology in programming environments, fairly rapid prototypes are possible. This is important because it permits us to make tentative decisions without incurring enormous commitments and to fine-tune the model by increasing its specificity.

An important effect of this method is that we learn from the program. One of the the most difficult parts of testing any complex theory is finding out whether the interaction of many subparts, each thought to be correct and well-understood, will in fact be harmonious, or whether it produces chaos when it is all put together. It is hard to imagine a better instrument for such a test than a computer. Since the process model is pointing out where the theory needs to be revised, we have a built-in guard against autocorrelation. Our previous work on generating stories by simulating the problem-solving behavior of a set of individuals [24] provides ample evidence of this effect.

Both principal investigators on this proposed project have extensive experience performing this kind of research. The programming will be done on UCI's DECsystem-10, which currently has 512K words of physical memory, virtual memory, and the UCI LISP system.

5.0 Bibliography

1] F. C. Bartlett.
Remembering: A Study in Experimental and Social Psychology.
Cambridge University Press, 1932.

2] Joseph Becker.
A model for encoding experiential information.
In Schank and Colby [36], 396-434.

3] Lawrence Birnbaum and Mallory Selfridge.
Problems in conceptual analysis of natural language.
Research Report 168, Yale Computer Science Department, 1979.

4] Gordon Bower.
Comprehending and recalling stories.
Presidential address, American Psychological Association, 1976.

5] J. Bransford and J. Franks.
The abstraction of linguistic ideas.
Cognitive Psychology 2:331-350.

6] Jaime G. Carbonell.
Subjective Understanding: Computer Models of Belief Systems.
PhD dissertation, Yale University, 1979.
Research Report 150, Yale Computer Science Department.

7] Eugene Charniak.
A framed PAINTING: the representation of a common sense knowledge fragment.
Cognitive Science 1(4):355-394, 1977.

8] B. N. Colby.
Computer assisted text ethnography.
School of Social Sciences, UC Irvine, February 1979.
Draft prepared for the UC Riverside Conference on Anthropology, Psychology, and Cognitive Structures.

9] B. N. Colby.
A partial grammar of Eskimo folktales.
American Anthropologist 75:645-662, 1973.

10] Allan Collins.
Processes in acquiring knowledge.
In R. C. Anderson, Rand J. Spiro, and W. E. Montague, editors, Schooling and the Acquisition of Knowledge, Lawrence Erlbaum, Hillsdale, New Jersey, 1976.

11] Allan Collins and M. Ross Quillian.
Retrieval time from semantic memory.
Journal of Verbal Learning and Verbal Behavior 8:240-247, 1969.

12] Richard E. Cullingford.
Script Application: Computer Understanding of Newspaper Stories.
PhD dissertation, Yale University, 1977.
Research Report 116, Yale Computer Science Department.

13] Edward A. Feigenbaum and Julian Feldman, editors.
Computers and Thought.
McGraw-Hill, New York, 1963.

14] Charles Fillmore.
The case for case.
In Emmon Bach and R. Harms, editors, _Universals in Linguistic Theory_, Holt, Rinehart, and Winston, New York, 1968.

15] Anatole V. Gershman.
Knowledge-Based Parsing.
PhD dissertation, Yale University, 1979.
Research Report 156, Yale Computer Science Department.

16] Richard H. Granger.
Adaptive Understanding: Correcting Erroneous Inferences.
PhD dissertation, Yale University, 1980.
Research Report 171, Yale Computer Science Department.

17] Richard H. Granger.
When expectation fails: towards a self-correcting inference system.
Proceedings of the First National Artificial Intelligence Conference, 1980 (to appear).

18] Walter Kintsch.
Memory for prose.
In Charles N. Cofer, editor, _The Structure of Human Memory_, W. H. Freeman, San Francisco, 1976, 90-113.

19] Walter Kintsch.
The Representation of Meaning in Memory.
Erlbaum Associates, Hillsdale, New Jersey, 1974.

20] Wendy G. Lehnert.
The Process of Question Answering.
Lawrence Erlbaum, Hillsdale, New Jersey, 1978.

21] Peter H. Lindsay and Donald A. Norman.
Human Information Processing.
Academic Press, New York, 1972.

22] Drew V. McDermott.
Assimilation of new information by a natural language-understanding system.
AI TR-291, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1974.

23] James R. Meehan.
Everything you always wanted to know about authority structures but were unable to represent.
Proceedings of the First National Conference on Artificial Intelligence, 1980 (to appear).

24] James R. Meehan.
The Metanovel: Writing Stories by Computer.
PhD dissertation, Yale University, 1976.
Garland Publishing, New York, 1979.

25] Marvin Misky, editor.
Semantic Information Processing.
MIT Press, Cambridge, Massachuetts, 1968.

26] Allen Newell and Herbert A. Simon.
Human Problem Solving.
Prentice-Hall, 1972.

27] Donald A. Norman.
Memory, knowledge, and the answering of questions.
CHIP Memo 25, Center for Human Information Processing, University of California, San Diego, 1972.

28] Donald A. Norman and David E. Rumelhart, editors.
Explorations in Cognition.
W. H. Freeman, San Francisco, 1975.

29] Jean Piaget.
The Language and Thought of the Child.
World Publishing, New York, 1955.

30] Christopher K. Riesbeck.
Conceptal analysis.
In Schank [34], 83-156.

31] Christopher K. Riesbeck and Roger C. Schank.
Comprehension by computer: expectation-based analysis of sentences in context.
Research Report 78, Yale Computer Science Department, 1976.

32] Charles J. Rieger.
Conceptual memory and inference.
In Schank [34], 157-288.

33] David E. Rumelhart, Peter H. Lindsay, and Donald A. Norman.
A process model for long-term memory.

In Tulving and Donaldson [40], 1972.

34] Roger C. Schank.
Conceptual Information Processing.
American Elsevier, New York, 1975.
Includes contributions by Neil M. Goldman, Charles J. Rieger, and Christopher K. Riesbeck.

35] Roger C. Schank and Robert P. Abelson.
Scripts, Plans, Goals, and Understanding.
Lawrence Erlbaum, Hillsdale, New Jersey, 1977.

36] Roger C. Schank and Kenneth M. Colby, editors.
Computer Models of Thought and Language.
W.H. Freeman, San Francisco, 1973.

37] Mallory G. Selfridge.
A Process Model of Language Acquisition.
PhD dissertation, Yale University.
Research Report 172, Yale Computer Science Department, 1980.

38] Oliver G. Selfridge.
Pandemonium: a paradigm for learning.
In Proceedings of the Symposium on Mechanism of Thought Processes, 511-529.
National Physical Laboratory, Teddington, England.
London, H. M. Stationery Office, 1959.

39] Perry W. Thorndyke.
Pattern-directed processing of knowledge from texts.
In Waterman and Hayes-Roth [41], 347-360.

40] E. Tulving and W. Donaldson, editors.
Organization of Memory.
Academic Press, New York, 1972.

41] D.A. Waterman and Frederic Hayes-Roth, editors.
Pattern-Directed Inference Systems.
Academic Press, New York, 1978.

42] Robert Wilensky.
Understanding Goal-Based Stories.
PhD dissertation, Yale University, 1978.
Research Report 140, Yale Computer Science Department.

43] Yorick Wilks.
An Artificial Intelligence approach to machine translation.
In Schank and Colby [36], 114-151.

44] Yorick Wilks.
Grammar, Meaning, and the Machine Analysis of Language.
Routledge and Kegan Paul, London, 1972.

[45] Yorick Wilks.
An intelligent analyzer and understander of English.
Communications of the ACM 18(5):264-274, 1975.

[46] Donald A. Norman and Daniel G. Bobrow.
On the role of active memory processes in perception and cognition.
In Charles N. Cofer, editor, *The Structure of Human Memory*, W. H. Freeman, San Francisco, 1976.