JUDGMENTAL INFERENCE:
A Theory of
Inferential Decision-Making During Understanding

Richard H. Granger, Jr.
Artificial Intelligence Project

Technical Report #182

April 1982

This research was supported in part by the National Science Foundation under grant IST-8120685, and by the Naval Ocean Systems Center under contract N00123-81-C-1078.

TABLE OF CONTENTS

Judgmental Inference:
A Theory of
Inferential Decision-Making During Understanding

ABSTRACT

In the course of understanding a text, a succession of decision points arise at which readers are faced with the task of choosing among alternative possible interpretations of what they're reading. Careful analysis of a wide range of sample texts reveals that such decisions are often based on complex evaluations of the interpretation being constructed, and sometimes cause the reader to construct and discard a number of intermediate inferences before settling on a final interpretation for a text.

This paper introduces Judgmental Inference theory as a proposed scheme of evaluation metrics and mechanisms, derived from examination of inference decisions arising during text understanding. A series of programs, ARTHUR, MACARTHUR and JUDGE are described, which incorporate some of the metrics and mechanisms of Judgmental Inference, enabling them to understand texts more complex than those that can be handled by other understanding systems.

1.0 Introduction

Many national newspapers carried front-page versions of the following story early this year:

> [1] A Nicaraguan soldier, who last year made a public statement alleging Cuban, Ethiopian and Nicaraguan military aid to Salvadorian leftist guerrillas, today publicly retracted his story at a State Department news conference.

Why did the Nicaraguan soldier make the statements he made, a year ago and now? Why did the State Department hold these two news conferences? It is possible that the State Department had some reason for holding the news conference, intending the Nicaraguan soldier to recant; but most readers assume that the State Department had different intentions that were not fulfilled, for reasons out of their control. Indeed, most readers don't even consciously think of the former interpretation, even though it is a logical possible alternative explanation of the events.

Texts describing complex situations like this often face readers with implicit choices among a number of alternative interpretations; and even apparently simple examples often present similar difficult decision points to an understander. Consider for instance:

This research was supported in part by the Naval Ocean Systems Center under contract N00123-81-C-1078.

[2] Kathy and Chris were playing golf. Kathy hit a shot into the rough. She wanted to let her good friend Chris win the game.

Most readers assume that the reason Kathy hit her shot into the rough was to increase her opponent's chances of winning, out of friendship. However, consider the following:

[3] Ken and Carl were playing golf. Ken hit a shot into the rough.

Most readers assume that Ken's shot was simply an accident, not intentional at all. However, this story is entirely identical to the first two sentences of [2] above. Hence, if a reader has made any inference at all by the time those two sentences of [2] have been read, then that inference is that Kathy's shot too was probably an accident. Yet by the end of the third sentence of [2], virtually no readers ever assume that interpretation.

Careful analysis of examples like these clearly indicate the existence of decision points at which human understanders are faced with the task of choosing particular inferential paths from among an array of possible alternatives. Our analysis shows that these inference decisions are based on complex evaluation metrics for judging the appropriateness of a particular inference in light of the context of previous inferences; textual and extra-textual factors influencing such decisions; and mechanisms for constructing and revising interpretations during understanding.

This paper introduces Judgmental Inference theory, which proposes a set of evaluation metrics and mechanisms derived from examination of inference decisions arising during text understanding. A series of programs are described which incorporate some of the metrics and mechanisms of judgmental inference: ARTHUR and MACARTHUR are working systems capable of understanding texts more complex than those that can be handled by other understanding systems, (e.g., PAM [Wilensky 1979], FRUMP [DeJong 1980], Ms.Malaprop [Charniak 1979], IPP [Lebowitz 1981]); and JUDGE, still under construction, incorporates a more integrated approach to judgmental inference than its two predecessors.

These theories are intended to be judged not only as working computer systems, but also as hypotheses about how people cope with inference decisions during understanding. Towards that end, some psychological and neurophysiological experiments are being designed to enable us to compare our theory's predictions with data on human understanders' inference behavior.

We view this work as compatible with and complementary to research that focuses primarily on representational issues in text understanding, such as Schank and Abelson [1977], Wilensky [1980], Charniak [1980]. By examining the occurrences of inference decisions during understanding, we intend to provide a look at the mechanisms by which such

representations are chosen, constructed, judged, confirmed and/or discarded during the processing of a text.

Following is an overview of the organization of the rest of this paper:

Section 2 presents analyses of some illustrative examples, thereby introducing some major categories of decisions a reader is faced with during understanding, and the evaluation metrics upon which such decisions are based.

Section 3 gives some examples of question-answering output from two working computer systems called ARTHUR and MACARTHUR, which use judgmental inference processes to successfully make inference decisions presented by these examples.

Section 4 dicusses additional metrics and mechanisms of Judgmental Inference which have been derived from more complex examples than the previous sections have shown. These are being incorporated into the design of the JUDGE program.

Section 5 summarizes the current state of the research so far, and speculates on some potential implications of this work for the areas of cognitive psychology and neurophysiology.

2.0 Illustrations of Understanders' Decisions

2.1 Evaluating and supplanting inferences

Recall the two "golf" stories given above:

[2] Kathy and Chris were playing golf. Kathy hit a shot into the rough. She wanted to let her good friend Chris win the game.

[3] Ken and Carl were playing golf. Ken hit a shot into the rough.

Let us examine in more detail the difficulties encountered by a reader of these examples. After reading the two-sentence version [3], readers infer that Ken and Carl both were playing to win, and that Ken's bad shot therefore was accidental, and will hinder his goal of winning the game. These inferences about Ken's intentions can be schematically illustrated as follows:

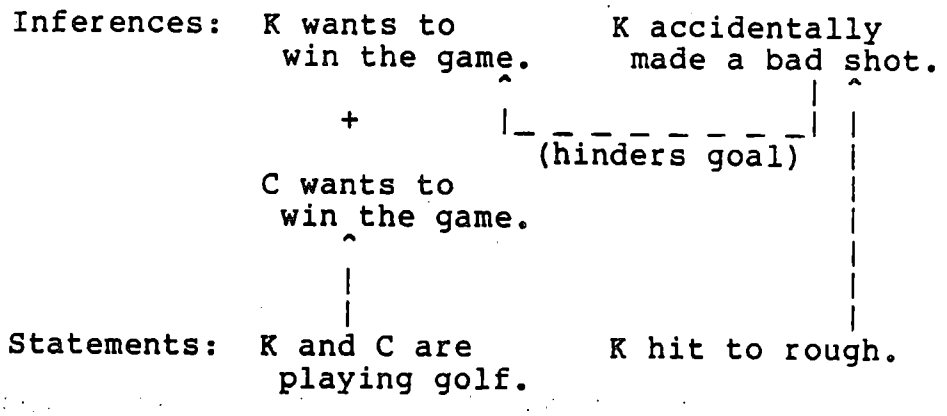

[Figure 1]

However, after readers have read the third sentence that appears in version [2], they appear to have changed this initial interpretation a great deal. It is not just that

Kathy doesn't want to win the game after all, but also that therefore she most probably made her bad shot on purpose, not accidentally. Virtually all readers arrive at this interpretation by the end of this example, by supplanting some of their initial inferences by new ones (see Granger [1980]).

2.2 The evaluation metrics of cohesion and parsimony

Why do people arrive at this different interpretation about Kathy's action in this example? The answer is far from obvious. In particular, there is no question of logical consistency here; once the initial inference about her goals has been supplanted then the interpretation is already entirely logically consistent. That interpretation, shown in Figure [2] below, would be that Kathy's bad shot was still simply an accident, even though it will help what we now know to be her goal of losing the game.

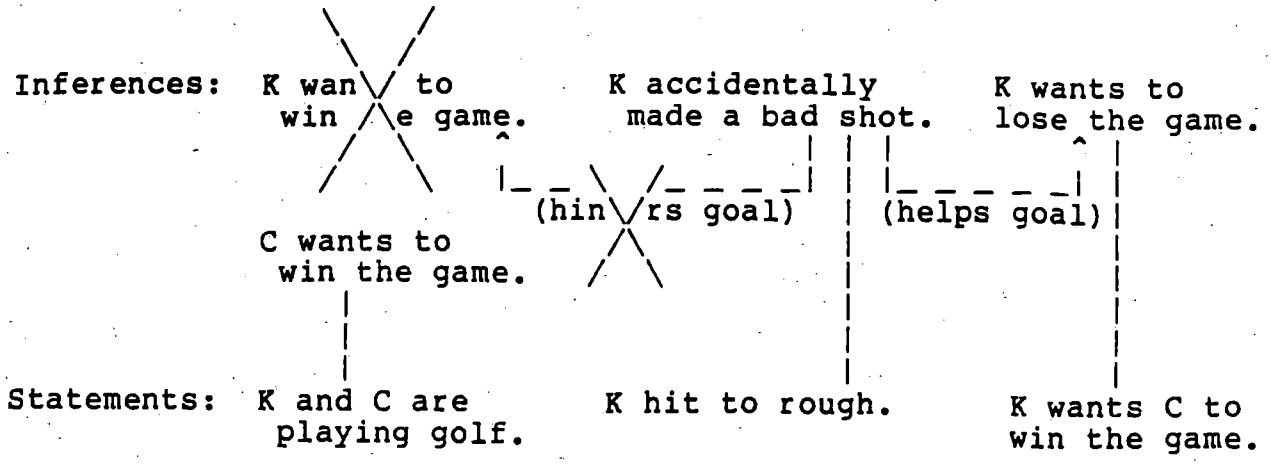

[Figure 2]

However virtually no readers ever arrive at this interpretation of this example; they all do the (unconscious) extra work of also supplanting the inference about K's bad shot; i.e., inferring that it was intentional, not accidental:

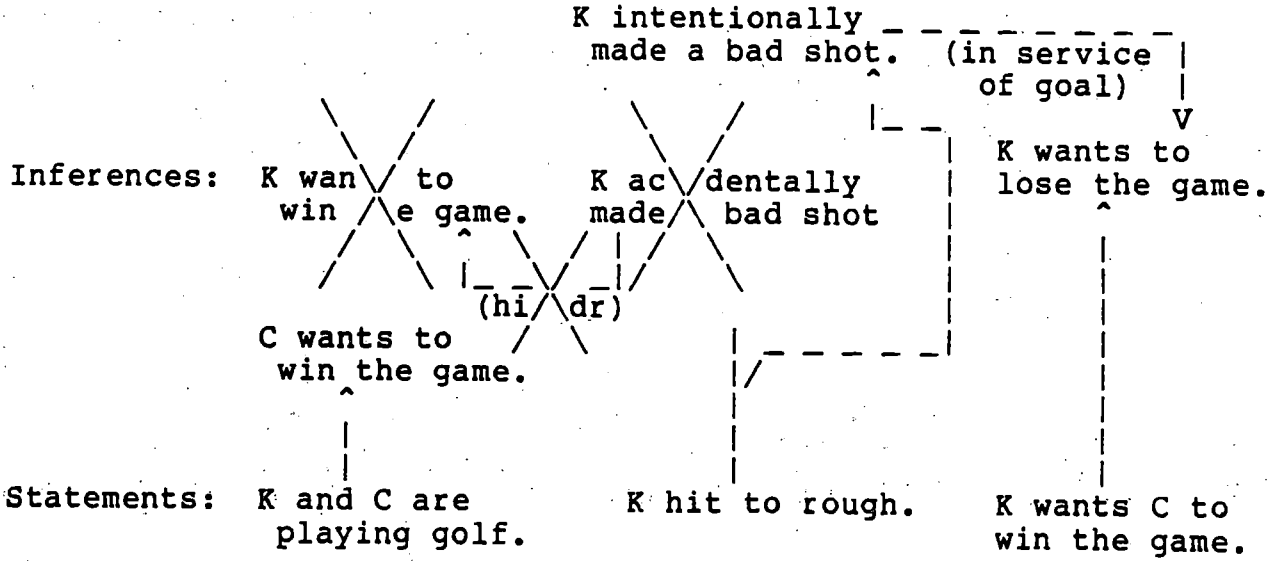

[Figure 3]

It turns out that the scope of this phenomenon is very wide: people often arrive at interpretations that appear to involve the supplanting of initial inferences, even when that extra work is not necessary on grounds of logical consistency. For instance, consider the following two examples:

[4] Doug went to a gas station. He robbed it and got away with $50.

[5] Mary picked up a magazine. She swatted a fly.

These are both examples, like [2] and [3] above, in which the most likely initial inferences (Doug was going to get gas for his car at the gas station; Mary was going to read the magazine) are supplanted by the time the reader has finished the example. (Again, this can easily be demonstrated by giving just the first sentences to readers and observing that they never infer that Doug might intend to rob the station, nor that Mary might intend to use the magazine as a fly swatter; a series of controlled protocol experiments along these lines is described in Rumelhart [1981]).

Again, these texts do not give rise to any logical contradictions, but they nonetheless imply initial inferences that human readers seem to eventually reject by the time they have finished reading the text. The decision to reject an initial inference, then, must depend on an <u>evaluation</u> of the representation based on some metric other

than logical consistency.

One such evaluation metric that was (implicitly) incorporated into previous theories of inference generation we have termed the "cohesion metric". Slightly differing versions of this measure of text representations have variously been referred to as "consistency" of interpretation (Rumelhart [1981]), "causal connectedness" (Schank [1973]), and "coherence" (Crothers [1978], Bower, Black and Turner [1979], Black [1980]). All of these hypotheses about text representations essentially say that the inferential net composing the representation must contain connections tying together all of the statements in text, via either referential, causal or intentional connective inferences. As Rumelhart [1981] puts it:

> "... the problem facing a comprehender is analogous to the problem that a detective faces when trying to solve a crime. In both cases there are a set of clues. The listener's (or reader's) job is to find a consistent interpretation of those clues." [p.30]

The cohesion metric therefore requires that every statement in a text be connected to at least one other, resulting in all the pieces of the text representation being tied together in a consistent fashion.

Cohesion by itself is not sufficient to evaluate the goodness of a text representation, however. Another evaluation metric, identified in our previous work (Granger [1980]), measures the parsimony of a representation, with

respect to the goals that motivate the events in the text. For instance, example [4] above could be interpreted in either of the following two ways:

(a) John went to the gas station intending to get gas, and then he changed his mind and decided to rob the station instead;

(b) John went to the gas station intending to rob it.

Both of these interpretations are not only logically consistent, but also referentially cohesive, since the actor (John) and the location (the gas station) were the same for both story statements under either interpretation. Furthermore, both interpretations are causally cohesive, since John had to get to the gas station before he could rob it, regardless of his intentions in performing those actions. Therefore, the cohesion metric does not differentiate between these two alternative interpretations, but people do: they universally seem to generate interpretation (b), and in fact they rarely even think of the possibility of (a) until it is pointed out to them.

Similarly, the following are all logically consistent and referentially and causally cohesive interpretations of example [5]:

(5a) Mary picked up a magazine to read it. She then was annoyed by a fly, and she swatted it with the magazine she was holding.

(5b) Mary picked up a magazine to read it. She then was annoyed by a fly, and she swatted it with a flyswatter that was handy.

(5c) Mary picked up a magazine to swat a fly with it.

The last interpretation (5c) reflects a story representation which consists of a single goal, getting rid of a fly, which both of Mary's actions were performed in service of. The other interpretations both consist of two separate goals, each of which explains one of Mary's actions. The evaluation metric of parsimony essentially tests that an interpretation be maximally parsimonious with respect to the number of goals used to explain the events in the story; i.e., the fewer separate motives inferred to account for the story events, the better.

2.3 Some related work, briefly contrasted

It is important to note that the evaluation metrics of cohesion and parsimony completely differentiate Judgmental Inference theory from the paradigm of "truth maintenance" (Doyle [1978, 1979], McDermott and Doyle [1980]). The theories behind truth-maintenance systems are concerned with the problem of maintaining a logically <u>self-consistent</u> database of propositions. Based on this notion, the various interpretations presented for the examples above (e.g., gas

station [4], fly swatting [5], golf [2],[3]) would all be considered equally self-consistent, since no explicit contradictions occur in any of these interpretations. The problem of course is that human readers universally select certain of these interpretations over the others, and in fact seem never to consciously notice other possible interpretations at all. The cohesion and parsimony metrics were developed to account for this observed preference.

The overall notion of making evaluations for goodness of fit and seeking alternatives when indicated by the evaluation is similar in spirit to the standard Artificial Intelligence notion of searching through a state-space, and such search techniques are covered in most AI textbooks. Those techniques do not apply to the problem of text understanding, however, for several reasons. The typical AI state-space represents a world of information that is simple, consistent, and correct (e.g., the domains of logical propositions or games such as chess, see Nilsson [1971]). The procedures for adding new information are well-defined theorems or rule-systems (e.g., production systems, see Newell and Simon [1972] or Winston [1977]). The evaluation metrics ("static evaluations") found in many search programs are numerical, making comparisons easy (see Barr and Feigenbaum [1981]). These programs focus on problem-solving skills and domain expertise. They seek alternatives by backing up, either in a timeless world (as in theorem-proving programs) or by reversing time (as in

many planning systems).

In contrast, the world of understanding actual text and human intentionality is complex and often inconsistent or incorrect (particularly where beliefs are concerned). First of all, time cannot simply be "reversed" in a cognitive model, the way it can in a mathematically-based problem-solving system: we cannot "unread" a piece of text. The strategies for seeking alternatives in our model must be psychologically more plausible than a mechanical scheme. Secondly, new information is added to an inference net by means of inference rules and inference contexts such as plans and goals, and hence the evaluation metrics described here are inherently non-numerical. Finally, the processes involved in understanding texts like those in this paper are dependent on knowledge of very different domains than those of traditional problem-solving systems, e.g., knowledge about social interactions, affective reactions to plan failures, and other instances of human intentionality in everyday situations.

2.4 The scope of the problem

As pointed out earlier, many texts ranging from apparently simple to complex give rise to complications in the understanding process, often requiring evaluation and re-evaluation of the constructed interpretation. So far, we have focused on an analysis of a few deceptively simple

examples, the "gas-station" story, "fly-swatting", and "golf". In general any similar "garden-path" text, i.e., one which gives rise to initial inferences that fail to be substantiated as the text continues, will present very similar processing difficulties to an understander. Following are a couple of additional examples; as before, to see the difficulty of processing them, imagine a question about <u>why</u> the actor did the first action, being asked initially after the first sentence and then again at the end of the example:

1. Collette picked up her sheet music. She began to fan herself.

2. Carl picked up the newspaper. He wanted to get his tennis raquet that was underneath it.

Additional (and more obvious) examples are those in which an understander expresses surprise at the ending of the example, having apparently consciously noticed the construction of the intermediate interpretation, that was subsequently supplanted. This subcategory of garden-path texts includes those that contain "surprise endings", jokes, and mysteries:

1. A man came to Ed's door, saying he wanted to talk to Ed about some insurance policies. Ed let the man in, and then the man pulled a knife and ordered Ed to give him all his money.

2. "I took my wife to the Bahamas, but then she found her way back home."

3. Willy North went into the Roger Sherman movie theater. He went up to the balcony, where Geoff was waiting. Geoff pulled out an ounce of cocaine and handed it to Willy. Willy paid him in hundred-dollar bills and left quickly.

Note that this last set of examples of garden-path texts tends to cause readers to express conscious surprise at the ending of the story, whereas the first set of texts seemed not to cause this kind of surprise. So far we have only tentative hypotheses to account for why some texts cause conscious surprise and others cause only this "unconscious surprise". We are working on a theory of "loose end" inferences that we hope will provide part of an answer, and we are designing some experiments that we hope will shed some light on the question.

2.5 Shaping interpretations of behavior

We have identified some further evaluations that understanders perform, beyond cohesion and parsimony, which arise when a reader is led to "doubt" any part of his interpretation of a text. Such doubts can be instilled either by information presented in the text, or by "extra-textual" factors (see Granger [1981]) which may steer the reader away from an otherwise plausible interpretation. Examples of such "doubt-factors" include the reader's knowledge of the reliability of the text source (e.g., the difference between the New York Times and the National

Enquirer); knowledge of an actor's deviousness (e.g., a car salesman vs. a priest); relative boredom or interest, i.e., the reader's desire to pursue possible alternative interpretations vs. just settling on a default interpretation that's "good enough". An easy way to induce a doubt factor in a reader is to simply tell him that his initial interpretation is incorrect; i.e., explicitly ask for a new and different interpretation of a text.

It turns out that readers are very capable of producing a series of such alternative interpretations of texts when they're continually told their initial interpretation is incorrect. For instance, following is a story adapted from a newspaper text, along with a series of interpretations informally elicited from a subject:

[6] The Pakistani ambassador to the United States made an unscheduled stop in Albania on his way home to what an aide of the ambassador described as "a working vacation".

Q1) Why did the ambassador go to Albania?
A1) It looks like he was on vacation -- he went to Albania first and then to home, I guess in Pakistan.

Q2) No, that's not the real reason. Why did he go to Albania?
A2) Well, maybe there was some emergency reason ... it said it was unscheduled, so maybe it was that something went wrong and they had to stop there, and then they went on.

Q3) Still not it, but try again; why did he go to Albania?
A3) Ok, maybe, well he's an ambassador, so he could have been supposed to go to Albania ... so it could have been a meeting, like "shuttle diplomacy" ... but it was supposed to be a secret, so that's why they said it was unscheduled.

These different interpretations of [6] are each based on different interpretations of the actor's reasons for doing what he did. It is natural that different behavior interpretations should give rise to different text interpretations; most current theories of text representation focus primarily on representation of the events described in the text, rather than on a more "syntactic" analysis of the structure of the text itself.

Our analysis of this and similar examples has revealed a large class of inference evaluations people perform based on their attempts to decide what kind of behavior an actor has performed, for instance:

1. "simple" goal pursuit, e.g., "John was hungry, so he ate a hamburger";
2. "complex" goal pursuit, (i.e., goal interactions; see Wilensky [1979]) -- e.g., "John wanted to see the football game but he also had a paper due the next day" (goal conflict);
3. deceptive or intentionally misleading behavior, e.g., "Clark wanted Lois to think he was drunk, so he smiled and fell off the barstool onto the ground";
4. accidental (non-goal-directed) behavior, e.g., "Jack smiled and fell off the barstool onto the ground"(!);
5. impromptu reactions to unplanned-for contingencies, e.g., "Bill threw himself under the jeep when he saw the man pull a gun".

(As suggested by some of the above examples, it will often be quite difficult or impossible for an understander to "correctly" place an actor's behavior into just one of these categories. This issue is discussed further in Section 4.)

We have devised a classification scheme for dividing up the gamut of possible interpretations of behavior, first dividing it into intentional vs. unintentional, and then further subdividing intentional behavior into simple, deceptive, pre-planned, impromptu, etc.; and unintentional or accidental behavior into various types of failures such as skill failure, information failure, etc.

Figure [4] illustrates the current state of our classification scheme for behaviors. We call each leaf of this tree an interpretation-"shape", since placing an actor's behavior at a particular leaf will result in a particular shape of the representation graph constructed, and because re-interpreting an actor's behavior results in re-shaping the representation. The next section will show some output of computer programs that use knowledge of cohesion, parsimony and interpretation shapes to construct interpretations of texts and to re-shape those interpretations when continually re-questioned about them.

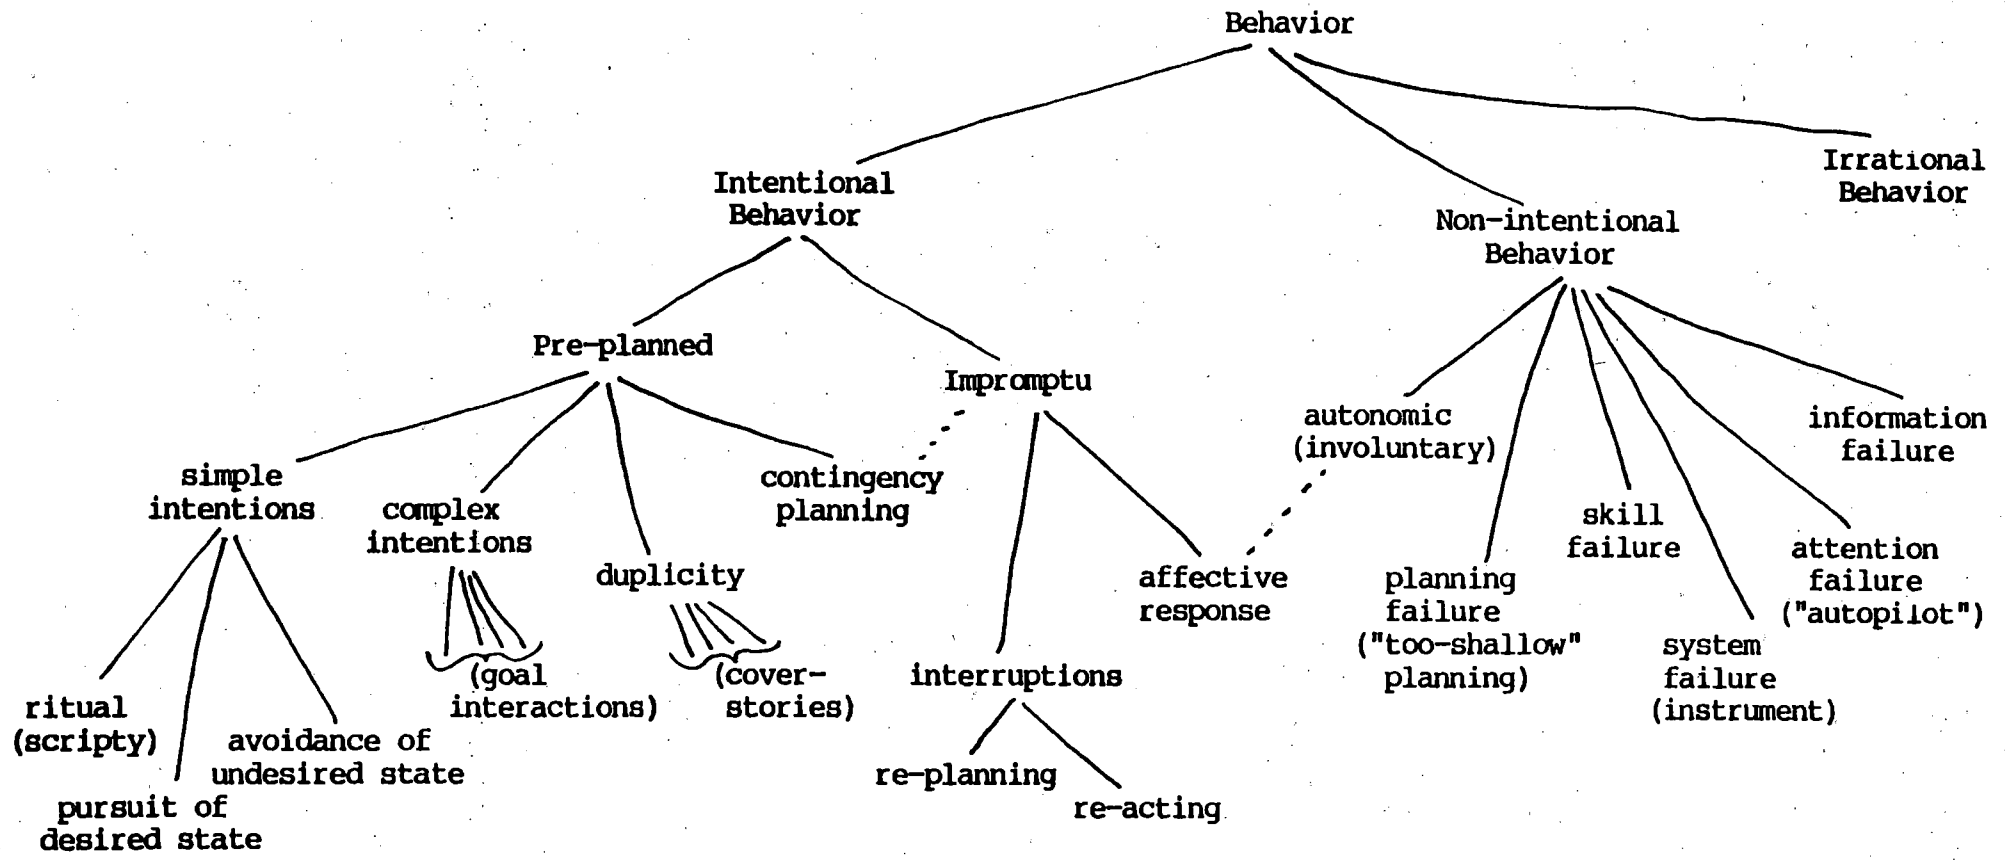
Behavior
Intentional Behavior
Irrational Behavior
Non-intentional Behavior
Pre-planned
Impromptu
simple intentions
complex intentions
contingency planning
duplicity
autonomic (involuntary)
information failure
skill failure
attention failure ("autopilot")
affective response
planning failure ("too-shallow" planning)
system failure (instrument)
ritual (scripty)
pursuit of desired state
avoidance of undesired state
(goal interactions)
(cover-stories)
interruptions
re-planning
re-acting

3.0 Judgmental Inference Systems in Action

3.1 ARTHUR changes its mind about its own inferences

ARTHUR uses the evaluation metrics of cohesion and parsimony, and the mechanisms of prediction and supplanting, to monitor its own progress through a text and arrive at the same interpretations of simple texts that people do. ARTHUR illustrates its understanding by answering questions about what it has read, including in its answers both its final interpretation and any intermediate interpretations that it had to discard (supplant) along the way. Following are the output from some sample question-answering sessions with ARTHUR on examples discussed above:

STORY1: Doug went to the gas station. He robbed it.

Q) Why did Doug go to the gas station?

A) AT FIRST I THOUGHT IT WAS BECAUSE HE WANTED TO GET SOME GASOLINE, BUT ACTUALLY IT'S BECAUSE HE WANTED TO GO THERE TO STEAL MONEY.

STORY2: Mary picked up a magazine. She swatted a fly.

Q) Why did Mary pick up the magazine?

A) 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.

STORY3: Kathy and Chris were playing golf. Kathy hit her ball into the rough. She wanted to let Chris win.

Q) Why did Kathy hit her ball into the rough?

A) AT FIRST I THOUGHT IT WAS ACCIDENTAL, BUT ACTUALLY IT WAS ON PURPOSE BECAUSE SHE WANTED TO LOSE THE GAME.

ARTHUR's representational scheme is based on that of Schank and Abelson [1977], including Conceptual Dependency representations for actions, script representations for stereotypical sequences of actions, and goal and plan representations for desired states and intended methods of achieving those states.

3.2 Annotated run-time output from ARTHUR

The following represents actual annotated run-time output of the ARTHUR program. The input to the program is the "swat fly" story given above:

[5] Mary picked up a magazine. She swatted a fly.

The first sentence causes ARTHUR to generate the plausible inference that Mary plans 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 an "explanation triple": a goal (being entertained), an event (picking up the magazine), and an inferential path connecting the event and goal (reading the magazine). The following ARTHUR output is generated from the processing of the second

sentence.

```
:CURRENT EXPLANATION-GRAPH:
 GOAL: (E-ENTERTAIN (PLANNER MARY) (OBJECT MAG))

 EV0: (GRASP (ACTOR MARY) (OBJECT MAG))
 PATH0: (READ (PLANNER MARY) (OBJECT MAG))
```

ARTHUR's explanation of the first sentence has a goal (being ENTERTAINed), an act (GRASPing a magazine) and an inferential path connecting the action and goal (READing the magazine).

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

The Conceptual Dependency for Mary's action: she struck a fly with an unknown object.

:FAILURE TO CONNECT TO EXISTING GOAL CONTEXT:

ARTHUR's initial goal inference (Mary planned to entertain herself by reading the magazine) fails to explain her action of swatting a fly.

```
:SUPPLANTING WITH NEW PLAUSIBLE GOAL CONTEXT:
 (PHYS-STATE (PLANNER MARY) (OBJECT MAG) (VAL -10))
```

ARTHUR now generates an alternative goal on the basis of Mary's new action: she may want to destroy the fly, i.e., want its physical state to be -10. This new goal also serves to explain her previous action (getting a magazine) as a precondition to the action of swatting the fly, once ARTHUR infers that the magazine was the INSTRument in Mary's plan to damage the fly.

```
:FINAL EXPLANATION-TRIPLE:
 GOAL1: (PHYS-STATE
         (PLANNER MARY) (OBJECT FLY) (VAL -10))

 EV1: (GRASP (ACTOR MARY) (OBJECT MAG))
 PATH1: (DELTA-CONTROL (PLANNER MARY) (OBJECT MAG)
 EV2: (PROPEL (ACTOR MARY) (OBJECT MAG) (TO FLY))
 PATH2: (CHANGE-PHYS-STATE (PLANNER MARY)
         (OBJECT FLY) (DIRECTION NEG) (INSTR MAG))
```

This representation says that Mary wanted to destroy a fly (GOAL1), so she planned to damage it (PATH2). Her first step in doing so was to get an instrumental object (PATH1). These two plans were realized (EVents 1,2) by her picking up a magazine and hitting the fly with it.

: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 the inferred goal under- lying Mary's action of GRASPing the magazine. This answer is generated according to ARTHUR's supplanted inference about the action (READ) and the active inference about the action (CHANGE-PHYS-STATE). The English generation mechanism used is described in Granger [1980].

3.3 MACARTHUR constructs alternative interpretations

MACARTHUR (Granger [1981]) has all of ARTHUR's inference abilities and more: it can construct a series of alternative interpretations of a text when it is led to "doubt" its initial inferences by being re-questioned. MACARTHUR uses knowledge of interpretation shapes described above to "re-shape" its initial interpretation of a text. In particular, MACARTHUR has knowledge suggesting a goal interpretation for a particular action in the context of a particular "shape". Table [5] illustrates some of the default goals that MACARTHUR will infer given particular actions and shapes.

	PURSUE-STATE	AVOID-STATE	COVER-STORY	NON-INTENTIONAL
(PTRANS (ACTOR Z) (OBJ Z) (TO X) (FROM Y))	Z pursuing function of LOC(X)	Z avoiding function of LOC(Y)	Z hiding secret mtg at LOC(X)	Motor failure: Z reacting to trouble with the INSTRUMENT (e.g. vehicle) of the PTRANS
(ATRANS (ACTOR Z) (OBJ X) (TO Y) (FROM Z))	Z wants Y to pursue function of OBJ(X)	Z avoiding function of OBJ(X)	Z hiding some other function of OBJ(X)	Information failure: Z thinks OBJ(X) is another object
(INGEST (ACTOR Z) (OBJ X) (TO Z's STOMACH))	Z pursuing function of OBJ(X)	Z avoiding neg effects of lack of OBJ(X)	Z hiding some other function of OBJ(X)	Information failure: Z thinks OBJ(X) is another object
(PROPEL (ACTOR Z) (OBJ X) (TO Y))	Z pursuing neg-phys-state of Y or of X	Z avoiding function of OBJ(X)	Z hiding goal of DPROX X to Y	Skill failure: Z wanted OBJ(X) to go to LOC(W)

Table 5: Shaping the Interpretation of Events

For each of these event/shape pairs, examples abound. For example, the pair <PTRANS/COVER-STORY> corresponds to the "secret meeting" explanation of story [6]. The pair <PROPEL/NON-INTENTIONAL> could arise if Z PROPELed X towards W, but it went to Y instead; (e.g., Ken hit a shot into the rough). Z's observed action of PROPELing X to Y might be inexplicable without reference to some possible "skill failure" on Z's part. Similarly, the pair <INGEST/NON-INTENTIONAL> could arise if Z INGESTed something that he wouldn't have INGESTed had he known what it was; e.g., he ate poison disguised as chocolate. One more: the pair <ATRANS/AVOID-STATE>: Z could have given X to Y not because he wanted Y to have it particularly, but because he (Z) wanted NOT to have X, because of some negative attribute; e.g., X is a TV and it distracts him from his work, so he wanted to get rid of it. The PURSUE-STATE explanations correspond mostly to likely default reasons for the event being explained, e.g. going somewhere (PTRANS) because you want to make use of some known function of the location, such as going to a singles bar to meet someone or going to a store to buy something. Following is some sample output from MACARTHUR after reading a simplified version of the above "Pakistani ambassador" story [6].

Q) Why did the ambassador go to Albania?
A) HE WENT ON A VACATION IN ALBANIA AND PAKISTAN.

Q) Try again. Why did he go to Albania?
A) MAYBE HE WANTED TO MEET WITH THE GOVERNMENT OF ALBANIA, BUT HE WANTED TO KEEP IT A SECRET.

3.4 Annotated run-time output from MACARTHUR

The following represents actual annotated run-time output of the MACARTHUR program. The input to the program is the Conceptual Dependency representation of the following story:

[7] Dr. Fitzsimmons yawned loudly. He left Carney and Samuelson and went into the next room. He opened the refrigerator.

MACARTHUR generates inferences connecting the first two statements, inferring that Fitzsimmons is tired and may be about to go to bed. The following MACARTHUR output is generated from the processing of the third conceptualization and subsequent questions.

```
:CURRENT EXPLANATION-GRAPH:
GOAL0: (S-SLEEP (PLANNER FITZ))
STATE0: (TIRED (ACTOR FITZ))
 EV0: (INGEST (ACTOR FITZ)
              (OBJECT AIR) (MANNER YAWN))
 PATH0: (INVOL-REACTION (ACTOR FITZ))
 EV1: (PTRANS (ACTOR FITZ) (OBJECT FITZ)
              (TO ROOM1) (FROM ROOM0))
 PC1: (D-PROX (PLANNER FITZ)
              (OBJECT FITZ) (TO ROOM1))
 PATH1: ($GO-TO-BED (PLANNER FITZ))
```

MACARTHUR's explanation of the first two statements consists of an "explanation graph" (Granger [1980a]), containing a goal (being rested) and two actions (yawn, change rooms), each of which are connected to the goal via an inferential path (an involuntary reaction, and the "go-to-bed" script, a known plan for get- ting rest). Now the third statement is read.

```
:NEXT SENTENCE CD:
(PTRANS (ACTOR FITZ) (OBJECT DOOR (PARTOF FRIDGE))
        (FROM CLOSED-POS) (TO OPEN-POS))
```

The Conceptual Dependency for Fitzsimmons' action; he moved the door to its open position.

```
:FOUND PLAN EXPLANATION:
(DRINK (PLANNER FITZ) (OBJECT MILK (TEMP WARM)))
CONTROL PRECONDITION IS:
(D-CONT (PLANNER FITZ) (OBJECT MILK) (FROM FRIDGE))
```

MACARTHUR connects a possible plan for going to bed (drinking warm milk) with the knowledge that milk is a food and therefore may be in the fridge. MACARTHUR now re-explains the previous action of changing rooms in terms of this plan of drinking warm milk.

PROXIMITY PRECONDTION IS:
(D-PROX (PLANNER FITZ) (OBJECT FITZ) (TO ROOM1))

:EXPLANATION COMPLETE:
:SHAPE IS PURSUE-DESIRED-STATE:

:READY FOR QUESTIONS:

Q) Why did Fitzsimmons open the refrigerator?

A) BECAUSE HE WANTED TO GET SOME MILK FROM IT SO HE COULD DRINK SOME WARM MILK SO HE COULD GET SOME REST.

MACARTHUR generates an answer to the question by using its current explanation-graph to find the goal for this action. It also mentions the intermediate inferences in the connective chain from event to goal. (MACARTHUR's English generation algorithm is taken from ARTHUR's, described in Granger [1980a].)

Now the same question will be asked again, causing MACARTHUR to re-evaluate its explanation.

Q) Are you sure?
Why did Fitzsimmons open the refrigerator?

:RE-PROCESSING AND RE-SHAPING EXPLANATION:
:NEW SHAPE IS AVOID-UNDESIRED-STATE:

(TIRED (ACTOR FITZ)) IS UNDESIRED STATE

:GENERATING NEW GOAL EXPLANATION:
GOAL: (P-AWAKE (PLANNER FITZ))

:ATTEMPTING TO SUPPLANT EXPLANATION:

MACARTHUR uses its classification of explanations into shapes (see section 3.2) to generate an alternative goal, that Fitzsimmons wants to stay awake, i.e., to avoid sleeping. It then attempts to re-explain the events of the story in terms of this new goal, thereby supplanting the previous goal explanation. This requires MACARTHUR to find connections between the story events and the new goal.

```
:FOUND PLAN EXPLANATION:
(DRINK (PLANNER FITZ) (OBJECT COKE))
CONTROL PRECONDITION IS:
(D-CONT (PLANNER FITZ) (OBJECT COKE) (FROM FRIDGE))

PROXIMITY PRECONDTION IS:
(D-PROX (PLANNER FITZ) (OBJECT FITZ) (TO ROOM1))
```

MACARTHUR has found a connective explanation: a known plan for avoiding sleep is to drink coke, and coke is food, so it may be found in the refrigerator, so going to the fridge and opening it can be interpreted as satisfying the preconditions of getting to the coke and getting hold of it in order to drink it.

```
:EXPLANATION COMPLETE:
:SHAPE IS AVOID-UNDESIRED-STATE:

:RE-RESPONDING TO PREVIOUS QUESTION:
Why did Fitzsimmons open the refrigerator?

A) AT FIRST I THOUGHT IT WAS BECAUSE HE WANTED TO
  GET SOME MILK FROM IT SO HE COULD DRINK SOME WARM
  MILK SO HE COULD GET SOME REST, BUT ACTUALLY IT
  MIGHT BE BECAUSE HE WANTED TO GET SOME COFFEE
  FROM IT SO HE COULD DRINK SOME COFFEE SO HE COULD
  STAY AWAKE.
```

MACARTHUR's final answer expresses both its supplanted initial explanation ("At first I thought ..."), that he might have intended to go to sleep because he was tired, along with its new explanation ("but actually ...") that his reaction to the state of being tired might have been to try to avoid going to sleep.

(See Granger [1981] and Schulenburg [1982] for a more complete discussion of MACARTHUR.)

4.0 Additional Categories of Inference Decisions

4.1 "Suspicious" understanding

It is often impossible for an understander to identify the "correct" interpretation shape for an actor's behavior. For instance, consider the following version of a story that was on the front page of a number of national newspapers earlier this year:

[8] A report by the New York State Racing and Wagering Board released today states unequivocally that leading jockeys conspired to "fix" at least 13 races in the mid-1970's, and that the jockeys have been "patently unbelievable" in denying their involvement in the scheme.

Understanding [8] requires the recognition that the observed behavior of jockeys can be very difficult to classify as either "accidental" or "deceptive". Hence, a jockey (or a jai-alai player, boxer, etc.) may lose a competition without an observer's being able to tell whether he did it intentionally or accidentally.

These are special cases of the general problem of detecting deceptive behavior by using knowledge of "cover stories". Some recent work in AI (e.g., Bruce and Newman [1978]) has pointed out that a method of maintaining separate "belief spaces" for different actors is crucial for understanding deception. However, understanding deception can also require a great deal more than this; in particular, a more subtle deceiver will typically try to

cause observers to infer for themselves some false interpretation of his actions, thereby covering up the real reasons. Political propaganda (such as falsely explaining the reasons for some military action), advertisements for products (falsely associating the use of a product with some desired outcome), and face-saving "white lies" (covering up potentially embarrassing reasons for your behavior) are all examples of this kind of deception. The ability to understand (and generate) complex deceptive behavior such as this depends not only on separate belief spaces, but also on the ability to construct plausible alterative explanations for events. The more plausible the alternative explanation, the more likely the deception is to succeed in misleading understanders.

A "suspicious" understander is one who can (at least) construct alternative interpretations of events, and then can attempt to decide among them, typically by gathering additional information. (Such information-gathering is based on finding a possible motive, i.e., finding a plausible explanation that the "obvious" explanation is intended to cover. The JUDGE program, currently under construction, is being designed to make use of knowledge of the shapes of alternative interpretations to detect plausible cover stories in the domain of criminal investigation.

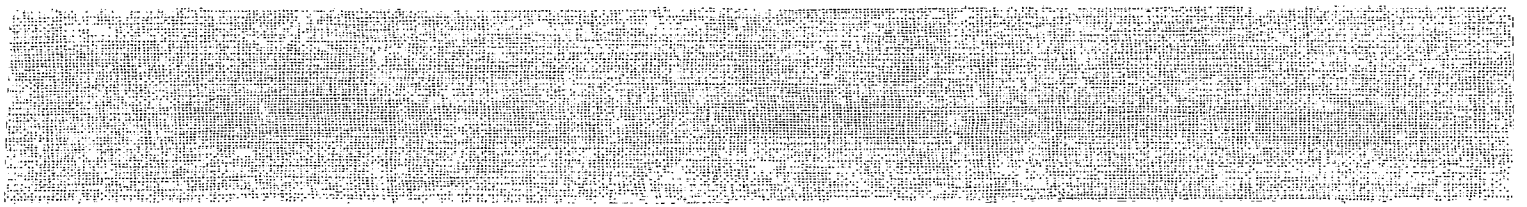

4.2 Cover stories

The theory of cover stories is based on two premises: first, that people often generate plausible default inferences during understanding, even though those inferences may turn out to be incorrect; and second, that presenting appropriate "clues" or facts to understanders can cause them to make false inferences, even when they're not directly lied to. For instance, if John rubs his stomach and licks his lips, and then goes to the kitchen, Mary (if she's watching) is very likely to make the appropriate inference that he's going there to get some food. However, if John actually has some other reason for going in there, say, to make a private phone call, he may want to cause Mary to make an inference about his hunger. To do this, he need only go through the above scenario, and Mary will infer he's hungry. Hence, John has fooled Mary by using his knowledge of typical default inferences to act out a scenario from which he knows an observer is most likely to infer a particular default conclusion, even when that conclusion is not the correct one.

The theory of cover stories deals with precisely this issue of how an actor can construct a scenario or make a statement in such a way that there are two or more possible interpretations of the actions or statements, and furthermore that the real motivation for the action, the one the actor wants hidden, is not the one that is the most

likely default interpretation. In this way a cover story can mislead a casual understander into inferring the wrong one of several alternative interpretations.

The other side of this issue is that an astute or suspicious understander can search harder through the space of possible interpretations, looking for possible hidden ones that he might otherwise have missed. At one extreme, this ability resembles paranoid behavior, with the understander constantly searching for hidden meanings underlying even the most innocent statements or actions. (Note that our intention here is not at all the same as Colby's [1973] model of paranoid behavior. Our intention is not to simulate a form of behavior, but rather to uncover the mechanisms, such as re-shaping interpretations, that might underlie such behavior in an understander.)

Putting this extreme case aside, however, it is clear that some understanding tasks require this type of suspicion by the reader: e.g., understanding political propaganda, and seeing through to why the statements were really issued; understanding advertisements, looking for the implications the advertiser inserts to mislead the understander without actually lying; attempting to solve a mystery such as a crime that has been committed, sifting through possibly misleading statements or testimony by suspects involved in the crime.

4.3 Natural subcategories of cover stories

Cover stories are composed of a "false" interpretation concealing the "true" interpretation or true motive for the behavior described. Both the false and true interpretations will have a particular "shape", as described earlier (e.g., pursue-desired-state, avoid-undesired-state, interruption-reaction, skill-failure accident, etc.) We have noticed that particular <u>combinations</u> of the shapes of the two interpretations comprising a cover story have certain identifiable natures. For instance, consider the following situation:

> Fred is a car salesman, trying to sell a car to Joe. Fred tells Joe that he feels friendly towards him, and that he's therefore giving him an especially good car for a better price than normal. Joe infers that Fred is simply hustling him.

Joe presumably possesses the societal knowledge that car salesmen may lie in order to sell cars. Hence he suspects that Fred may not be telling the truth about either the merits of this car, nor about his apparent friendly feelings for Joe. Rather, Fred may be constructing a cover story (i.e., that he wants to do Joe a good turn) to hide his simple pursuit goal of making a commission from the sale of a car.

Contrast this with the following example:

> Prof. Cooper tells his colleague Bernstein that he is looking forward to relaxing this summer and getting a much-needed rest. Bernstein immediately infers that Cooper didn't get the research grant that he knew Cooper had applied for.

In this situation Cooper has evidently constructed a cover story in which he appears to be pursuing a vacation, whereas actually he is simply re-planning his summer as a reaction to the unplanned-for fact that he now has no summer salary coming in.

We can differentiate these two cover stories in terms of the combinations of shapes comprising the cover story. In the first situation, Fred the salesman is using a pursuit-based cover story (wanting to do a favor for his friend Joe) to conceal a pursuit-based underlying motive, (making a commission on the sale). This type of cover story consists of "pursuit covering pursuit". Similarly, one interpretation of the "Pakistani ambassador" story given earlier was that the ambassador made use of a pursuit-based cover story of wanting a vacation to conceal a pursuit-based underlying motive of meeting with the Albanian government. It seems that the typical motivation behind such cover stories can be characterized by "premeditated deception" by the speaker, in service of a hidden web of goals.

In the second situation given above, Cooper uses a pursuit-based cover story of wanting a vacation to conceal his interruption-reaction to the presumably unplanned circumstance of not having a salary. Consider another

example of this category of cover-story situation:

> Mary sees a mouse in her kitchen and promptly hops up onto a chair. "I wasn't afraid of that silly mouse", she claims, "I just wanted to reach up to clean off this spot on the ceiling".

Again, her pronouncement has the shape of pursuit (cleaning the spot) covering an interruption-reaction to the unplanned circumstance of seeing the mouse. These cover stories seem to have the property of "saving face" for the speaker, i.e., keeping the speaker from the embarrassment of admitting to the actual reason for his actions.

Other permutations of simple shapes combined to create cover stories give rise to subtly different categories of motivations for an actor to create a cover story.

We hope this line of research will lead us to a more complete theory of how understanders reconstruct interpretations of episodes, and how a storyteller or reporter can construct cover stories to conceal the real intentions underlying an episode. This research has possible applications in understanding and analyzing potentially misleading texts such as political reports and advertising messages, as well as in making deeper and more discriminating readings of other non-straightforward texts, and in constructing interpretations of observed episodes.

4.4 Understanding accidents

In the attempt to attribute intentions to actors on the basis of their actions, an understander will often confront an action which seems inexplicable when regarded as intentional behavior. Sometimes (as discussed in the previous section) this can be a clue that the actor has been acting out a cover story, in which case re-interpreting his previous actions in light of this new information may lead to discovery of his true motives. However, sometimes actions are inexplicable in terms of intentionality because they are either accidental or are immediate (and poorly planned) responses to some unforeseen contingency. Our initial investigation has led us to hypothesize four categories of accidents that can occur in an actor's plan. We call these classes skill failures, information failures, system failures and planning failures. The following discussion gives examples of all four of these accident types.

Recall the following:

> Ken and Carl were playing golf. Ken hit a shot into the rough.

As discussed earlier, if a reader is asked the question "Why did Ken hit a shot into the rough?" after reading this example, the most likely answer will be that he didn't do it on purpose at all; it was an accident. If we analyze this further, we see that Ken's action was probably intentional:

he did aim the club at the ball and swing in such a way as to apply force to the ball. Hence, he must have had a goal in so doing: we infer that it was to get the ball closer to the hole, so as to get there in the fewest possible strokes, in hopes of getting a good score in the game. However, the causal result of her intentional action did not turn out the way she intended: the ball went in a different direction than we infer that she wanted it to go. Recognizing possible skill failures such as this one is a necessary part of understanding actions by intentional actors.

Other types of failures can occur in the performance of an action intended to have a particular causal result; for example, consider:

[9] John was looking for some cereal, and he accidentally ate a bowl of rat poison flakes instead. He got very very sick.

We infer from this example that John didn't intend to get sick, but he did, of course, intentionally perform the actions involved in eating the rat poison. The problem here is that John didn't know the fact that the stuff he was eating was poison; the accidental outcome of this action was due to an information failure.

The third class of accidents we have looked at can be illustrated by the following example:

> John was driving up to Oxnard to visit his girlfriend. He got a flat tire when he was about halfway there.

This accident is classified as a "system failure", i.e., a problem with some instrument the planner is counting on to achieve his plan. Other examples of this category of accident include problems with the postal service, damaged tools, illness, etc.

Finally, the fourth class of accident we have identified are due to errors that occur at the time an actor is planning a future episode, not at the time of execution of the plan. These accidents still go unnoticed until the time of execution, however. For instance:

> Jack went to the theater to see "Empire". He realized when he got there that he had forgotten his wallet.

Jack's problem is clearly attributable to his own "too-shallow" planning at the time he left his house to go to the movie theater. He didn't bother to think ahead to the preconditions of the plan he was about to carry out, and hence he forgot to fulfill one of the simple preconditions, namely, having money to pay for his ticket.

Understanding stories involving failures of skill, information, an instrumental system or sufficient pre-planning requires the reader to acknowledge the possibility that an actor's plan might fail. Such a failure would be represented as a mismatch between the causal result of an action and the inferred intention in service of which

the action was performed. Hence, understanding such texts requires monitoring possible alternate causal results of actions, in addition to the possible alternate intentions in service of which the actions might have been performed.

4.5 Re-planning and re-acting

When an accident occurs, a planner will usually respond by re-planning around the accident and re-acting on the basis of the new plan. While there exist many planning and problem-solving systems capable of re-planning, we know of no systems that can detect or understand accounts of episodes involving re-planning. For example, recall the following situation:

> John was driving up to Oxnard to visit his girlfriend. He got a flat tire when he was about halfway there.

This accident is classified as a "system failure", as described earlier. John's possible responses include the following:

1. Resume-plan: John may decide to try to continue his current plan of driving up to his girlfriend's. For instance, he may simply keep driving (incidentally ruining his tire rim).
2. Local-remedy: John may decide to interrupt his plan, to attend to fixing the problem. E.g., he may replace the tire with a spare, and then resume his drive.
3. Change-plan: John may decide to abandon his existing plan and try another one to achieve the goal. E.g., he may try to hitchhike the rest of the way, or rent another car, or walk.

4. Abandon goal: John may decide to give up, and simply sit by the side of the road waiting for the police to drive by and rescue him. (Issues in understanding goal abandonment are discussed in Schank and Abelson [1977].)

In addition to these issues revolving around his existing plan, John may also have any of the following responses to the accident:

1. Diagnose: John may try to dig deeper into what went wrong with his tire.
    1. Was it simply worn down? If so, he may make a mental note to start having regular maintenance checks on his car.
    2. Was there something sharp on the road? If so, he may try to get rid of the object so that the next driver's tire won't be punctured by it.
    3. Did someone intentionally put a nail in the road? If so, he may want to find out who and why, to prevent the person from doing other bad things to him, and perhaps to exact revenge.
2. Emote: John may sit down in the road and cry, or get angry and kick the car, etc. Such affective reactions have been discussed by Lehnert [1980].

5.0 Conclusions: Where we've been, where we're heading

5.1 What we're proposing

We have observed that people's understanding behavior is marked by an ongoing process of making inference decisions. Among the decisions understanders implicitly make are:

1. Is the interpretation referentially and causally cohesive?

2. Is the interpretation parsimonious with respect to the actors' intentions?

3. Is there reason to doubt or be suspicious of the shape of the initial interpretation?

4. Is there reason enough to revise the interpretation (supplant, re-shape, etc) or should it be left with "loose ends"?

The evaluation metrics and the construction and revision processes of Judgmental Inference theory are derived directly from our observations and analyses of some of the classes of inference decisions that readers are faced with during the task of text understanding.

We view these theories as compatible with and complementary to theories of text representations, since we intend to describe the mechanisms by which such representations are chosen, constructed, judged, confirmed and/or discarded in the process of understanding. Our theories have so far been incorporated into two working computer programs, ARTHUR and MACARTHUR, and are currently being used as the design impetus for a new computer system called JUDGE, and for a series of psychological and neurophysiological experiments, briefly described below, to test the correspondence of our theories to people's actual understanding behavior.

5.2 Minds, brains and processes

A number of researchers in the neurosciences (e.g., Arbib [1979], Geschwind [1980]) have pointed out that brain research might help guide parts of cognitive science and AI research, and vice versa. One particular issue that has been pointed out frequently is that "there is no evidence for the existence of any all-purpose computer [in the brain]. Instead, there seems to be a multiplicity of systems for highly special tasks." (Geschwind [1980], p.191). Our research on inference decisions has indeed led us away from viewing human understanding behavior as arising from a "general purpose computer"; we have ended up instead deriving a number of special-purpose mechanisms, e.g., inference pursuit, evaluation, supplanting, re-shaping, which comprise our judgmental inference model of understanding.

We are currently designing a number of psychological and neurological experiments on inference decisions, based on the predictions of our model. We hope to use the data from such experiments to find problems with our theories, and to refine the model, thereby working eventually towards some small amount of "neurological validity" in our process models of cognition.

One of the major problems that our theories still have not adequately addressed is that of differentiating among inference decisions that people make unconsciously vs.

those that are consciously noticed. Furthermore, our initial investigation has shown that there are no simple rules for conscious vs unconscious inference decisions; in particular, there appear to be a number of individual differences among readers -- one view is that some readers consciously notice these decisions more easily than others; another view is that the readers who do not notice these decisions are not really making them -- rather they are postponing the decisions by leaving "loose ends" in their interpretations, and resolving them later on. Hence, our initial experiments are being designed specifically in an attempt to shed some light on this problem of conscious vs unconscious inference decisions. There are also a number of existing studies on this and related issues in inference processes during understanding (e.g., Rumelhart [1981], Crothers [1978], Hillyard and Kutas [1980], Black [1981]). We hope to re-interpret some of the data from these experiments in terms of the predictions of our model.

6.0 References

Arbib, M.A. and Caplan, D. Neurolinguistics must be computational. Behavioral and Brain Sciences, 2:449-483, 1979.

Black, John. The effects of reading purpose on memory for text. Cognitive Science Technical Report 7, Yale University, 1980.

Bower, Gordon H., John B. Black, and T. J. Turner. Scripts in text comprehension and memory. Cognitive Psychology 11:177-220, 1979.

Bruce, Bertram and Denis Newman. Interacting plans. Cognitive Science 2:195-233, 1978.

Charniak, Eugene. On the use of framed knowledge in language comprehension. Yale Computer Science Research Report 137, 1978.

Crothers, Edward J. Inference and Coherence. Discourse Processes, 1:51-71, 1978.

DeJong, G.F. Skimming Stories in Real Time: An Experiment in Integrated Understanding. Yale Computer Science Research Report 158, 1979.

Doyle, J. A truth maintenance system. Artificial Intelligence 12(3), 1979.

Geschwind, N. Neurological Knowledge and Complex Behaviors. Cognitive Science, 4:185-193, 1980.

Granger, R.H. Adaptive Understanding: Correcting Erroneous Inferences. Computer Science Research Report 171, Yale University, 1980a.

Granger, R.H. When expectation fails: Towards a self-correcting inference system. In Proceedings of the First National Conference on Artificial Intelligence, Stanford University, 1980b.

Granger, R.H. Directing and re-directing inference pursuit: Extra-textual influences on text interpretation. In Proceedings of the Seventh International Joint Conference on Artificial Intelligence (IJCAI), Vancouver, British Columbia, 1981.

Granger, R.H. Judgmental Inference: Inferential Decision-Making During Understanding. Proceedings of the Fourth Annual Meeting of the Cognitive Science Society, Ann Arbor, Michigan, 1982.

Granger, R.H. and K. Eiselt. 'Suspicious' Understanding: Detecting Possible Deception by Inferring Alternative Explanations. Computer Science Technical Report 185, University of California, Irvine, 1982.

Kutas, M. and S.A. Hillyard. Reading Senseless Sentences: Brain Potentials Reflect Semantic Incongruity. Science, 207:203-5, 11 Jan 1980.

Lebowitz, M. Generalization and Memory in an Integrated Understanding System. Computer Scinece Research Report 186, Yale University, 1981.

McDermott, D.V. and Jon Doyle. Non-monotonic logic I. Artificial Intelligence 13(1,2):41-72, 1980.

Meehan, J.R. Boy meets goal, boy loses goal, boy gets goal: the nature of feedback between goal-based simulation and understanding systems. Computer Science Technical Report 170, University of California, Irvine, 1981.

Rumelhart, D.E. Understanding understanding. Technical Report 100, Center for Human Information Processing, University of California, San Diego, January 1981.

Schank, R.C. Causality and Reasoning. Technical Report 1, Istituto per gli studi Semantici e Cognitivi, Castagnola, Switzerland, 1973.

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

Schulenburg, D. Interpreting Intentional and Unintentional Behavior. Technical Report, University of California, Irvine, 1982.

Wilensky, R. Understanding Goal-Based Stories. Research Report 140, Yale Computer Science Department, 1978. Garland Press, New York, 1980.