

STRATEGIST: A Program that Models  
Strategy-Driven and Content-Driven Inference Behavior

Richard H. Granger  
Kurt P. Eiselt  
Jennifer K. Holbrook

Artificial Intelligence Project  
Computer Science Department  
University of California  
Irvine, California 92717

ABSTRACT

In the course of understanding a text, different readers use different inference strategies to guide their choice of interpretations of the events in the text. This is in contrast to previous computer models of understanding, which all use the same (single) strategy while concentrating on details of content-driven inference. The separate strategies are theorized to be composed of the same component inference processes, but of different rules for application of the processes. The use of different strategies occasionally results in different interpretations of a single text. This paper presents both the results of new experimental data and a working computer program, called STRATEGIST, that models both strategy-driven and content-driven inference behavior. The rules which make up two of these strategies are presented.

1.0 Introduction

Artificial Intelligence models of human understanding have implicitly assumed a single strategy of inference behavior. The integrated understander's strategy usually goes:

1. While reading a sentence, make as many inferences as possible.
2. Connect inferences from the two sentences.

However, we have observed that not all readers make interpretations of text which conform to this strategy. For example, subjects in our recent experiments (Granger & Holbrook, 1983) read the following story:

[1] Nancy went to see a romantic movie. She was depressed.

Our experiments show that different individuals who read this story had two significantly different interpretations: (1) Nancy went to

-----  
This research was supported in part by the National Science Foundation under grant IST-81-20685 and by the Naval Ocean Systems Center under contract N00123-81-C-1078.

the movie to be entertained, and the movie depressed her (perhaps because it was romantic); vs. (2) something before the movie depressed Nancy, and she went to the movie to cheer up. Our experiments have indicated that at least two different inference strategies for interpreting text exist. However, these strategies are so closely related that, most of the time, readers using different strategies will come up with the same interpretation of the events related in the text.

We theorize that the same component inference processes which comprise each inference strategy are available to all readers. The difference in the strategies lies in the different rules used to apply the component processes. This paper presents these theorized processes and rules in a prototype model, called STRATEGIST, which exhibits the observed behavior of human readers.

Several researchers (e.g. Schank & Abelson, 1977) have hypothesized inference processes which allow the reader to interpret text. Psychological experiments such as those conducted by Graesser (1981), Rumelhart (1981), and Seifert, Robertson, and Black (1982) have determined when in the understanding process various types of inferences are made. However, these experiments were not designed to study the differences in processes which our experiments have discovered. The results of many of these studies can be reinterpreted in light of our results. The programs which were written to emulate human inference behavior (e.g. Wilensky, 1978; Granger, 1980; Wilensky, 1983) have also failed to model this particular aspect of inference decisions.

2.0 Background

Many story understanding systems have been written which can easily interpret simple text. Recall the example story:

[1] Nancy went to see a romantic movie. She was depressed.

PAM (Wilensky, 1978) would interpret this story by assuming that Nancy has the goal of entertaining herself. Going to the movie is

Nancy's plan for satisfying this goal. PAM would try to fit Nancy's depressed state into the plan that was executed, probably by inferring that the movie was depressing, or that the movie was not entertaining, and that she was depressed because she could not fulfill her goal. BORIS (Dyer, 1980) would come up with the same interpretation for the same reasons, even though it employs more complex knowledge structures. MACARTHUR (Granger, 1981) would be able to come up with both interpretations, but would always generate the same initial interpretation as the other systems.

These systems have all worked from a basic set of premises which include two types of rules. Content-driven rules are rules which generate inferences on the basis of the understander's specific knowledge of the situation described in the text. Strategy-driven rules are rules which generate inferences or suppress content-driven inferences using extra-textual considerations. In other words, the strategy-driven inferences themselves define the specific context of the situation described in the text. Humans understand stories using an Inference Manager which applies the content-driven rules, as well as the strategy-driven rules specific to their behavior. The kinds of inferences generated by both types of rules include explanatory inferences. These inferences explain why the stated events occurred. In other words, explanatory inferences are adding to the context. (For example, goals can be explanatory inferences with respect to intentional actions.) If explanatory inferences add enough context, they can give rise to predictive inferences, expectations about the events which will occur in the text. Plans are examples of predictions from goals. (Predictive inferences always 'look ahead' to account for some new input in the text; reciprocally, postdictive inferences are those plan inferences that look backward from an explanatory goal inference to account for previous events in the text.)

The following set of rules is content-driven, and used by all readers to understand text:

- C1. As a sentence is parsed, try to fit new input/conceptualizations into existing context.
- C2. If inferences conflict with specific statements in the text, the specific statements rule out the inferences, which are supplanted by interpretations which do not conflict with the specific statements.

All readers verify understanding by satisfying evaluation metrics, which the Inference Manager applies to the interpretation. There are at least two such

evaluation metrics:

- M1. All events should be causally related to each other (Cohesion).
- M2. Make the least complex interpretation of events possible (Parsimony -- Granger, 1980).

Our experiments have found that many subjects will indeed interpret story [1] as the systems described earlier do. We call people who come up with this interpretation Perseverers. Our data indicate that Perseverers will make inferences as soon as possible when reading text. These early inferences are the context in which further events are interpreted. Such readers persevere with an inference until a contradictory event or concept forces a change of interpretation. These are the Perseverer's strategy-driven rules:

- PS1. If there is no previous context, make default inferences.
- PS2. Inferences should always be as specific as possible. (Wilensky, 1983).

The Perseverer's set of strategy-driven rules is used by all of the systems discussed above.

Applying these rules, this is how we hypothesize that a Perseverer would go about interpreting story [1]:

INPUT: Nancy went to see a romantic movie.

Application of C1: This is the first sentence, so there is no previous context to constrain inferences. Several low-level inferences, basically an unstated part of the sentence, are made (e.g. Nancy not only went to the movie theater, but saw the movie as well).

Application of PS1 and PS2: These default inferences which are made are as specific as possible. The most important for our purposes is the explanatory inference that generates a goal to explain why Nancy went to the movie. This default goal is that Nancy wanted to be entertained by the movie. This explanatory inference gives rise to several predictive inferences, among them, the expectation that Nancy was happier after she saw the movie.

INPUT: She was depressed.

Application of PS1: Because there is previous context, C1 applies. The reader tries to fit this new sentence into the existing

context which the predictive inferences set up (Nancy was happier after she saw the movie).

Application of C2: The predictive inference that Nancy was happier must be supplanted with the specific knowledge that Nancy was depressed. The explanatory inference need not be supplanted, but the reader must realize that Nancy's goal of happiness was not fulfilled.

Application of M1 and M2: The most parsimonious explanation of the story is that Nancy's goal of happiness was not fulfilled because her goal of entertainment was not fulfilled. This explanation is also cohesive.

Note that M1 is constantly being applied to the two sentences, as connections are searched for. Notice also that there are other interpretations which do not assume that the goal of entertainment wasn't fulfilled -- for example, Nancy may have enjoyed the movie so much that she was depressed because it ended. However, this is not the most parsimonious interpretation.

There is a different initial interpretation which subjects in our experiments made, which is as plausible as the interpretation made by the Perseverers and all of the systems mentioned above. This interpretation is that Nancy was depressed before she saw the movie, and went to the movie to cheer up. We call people who make this interpretation Recencies. Recencies are readers who delay making inferences until enough information is present. A basic rule which drives this strategy is: when more text is available, and the text is ambiguous, leave a loose end (Granger, 1980), because later text will explain earlier events. The most recent inference will then become the context in which the earlier text is interpreted. To arrive at this alternate interpretation, a Recency must have a different set of strategy-driven rules from that of the Perseverer, which the Inference Manager applies:

RS1. If there is no active context, only low-level goals are to be inferred.

RS2. If there is no more text, inferences should be as specific as possible.

RS3. If there is more text, leave a loose end.

Applying these rules, this is how we hypothesize that a Recency would process story [1]:

INPUT: Nancy went to see a romantic movie.

Application of C1: This is the first sentence, so there is no context to direct inferencing.

Application of RS1 and RS3: Only low-level inferences are made (e.g. Nancy saw the movie). There is more text, so inferences are left as unspecified as possible, and loose ends are left rather than generating explanatory and predictive inferences.

INPUT: She was depressed.

Application of C1: The only existing context is low-level.

Application of RS2: There is no more text, so specific explanatory and predictive inferences must be made from the present concept. The most important explanatory inference for our purposes is that Nancy has a goal of alleviating her depression, and the predictive inference which follows is that Nancy will do something to alleviate her depression.

Application of C2: The causal relation is ambiguous, so the goal-based predictive inference (that Nancy will do something to alleviate her depression) is maintained, and a search is conducted for something in the text which can serve as this plan. Going to the movie fulfills this predictive inference, and so the final interpretation of the text is that Nancy went to the movie to cheer up.

Application of M1 and M2: The events are related to each other (Causal Cohesion satisfied). The fewest number of inferences were used to relate the events to one another (Parsimony satisfied).

(Although the explanation of rule application appears to have applied the rules in a particular order, we do not have any theories about rule ordering. The rules are probably applied as they become appropriate, but there is not necessarily a linear order for application.)

These two different strategies may seem to describe different sets of component processes altogether. This is deceptive; the processes are strikingly similar. The evaluation metrics which both strategies use are the same; Causal Cohesion must be satisfied for both strategies, although the cause/effect chains are different. Furthermore, both interpretations are parsimonious. There is no evidence that either strategy cannot make particular types of inferences. Both strategies make necessary inferences, such as connecting referents, inferring that Nancy went inside the theater, bought the ticket, and saw the movie (see Seifert, et. al, 1982, for discussion). We theorize that both strategies also make use of the same knowledge representations. Both strategies generate explanatory inferences and predictive inferences. With both interpretations, the inferences that are made affect the interpretation of other events.

It is the strategy-driven mechanism that drives the ongoing decision to either apply or suppress particular content-based inferences during understanding. Thus, on a given text, the same (potential) content-based inferences will be available regardless of strategy, but, depending on the strategy used, some of those "available" inferences will be generated while others will not. STRATEGIST's Inference Manager can apply the text-interpretation rules of either strategy, and so can derive either interpretation of an ambiguous text.

Following is a brief summary of the steps Perseverers and Recencies take during the processing of both the Nancy text and its reverse:

Text 1F (Forwards):

Nancy went to see a romantic movie.  
She was depressed.

Perseverer understanding steps

1. Explanatory inference of goal (be entertained)
2. Predictive inferences from goal (see-movie plan will succeed in satisfying entertainment goal)
3. Unsuccessful search for connection between 'plan success' postdiction and 'depression' affect
4. Successful search for connection between alternate 'plan failure' postdiction and 'depression' affect

Recency understanding steps

1. Leave loose end from first sentence
2. Explanatory inference of goal from 2nd sentence (alleviate depression)
3. Postdictive inference from goal (plan for alleviate-depression)
4. Successful search for connection between plan and event (see movie)

Text 1B (Backwards):

Nancy was depressed.  
She went to see a romantic movie.

Perseverer understanding steps

1. Explanatory inference of goal (alleviate depression)
2. Predictive inferences from goal (plan for alleviate-depression)
3. Successful search for connection between plan and event (see movie)

Recency understanding steps

1. Explanatory inference of goal from 2nd sentence (be entertained)
2. Postdictive inferences from goal (see-movie plan will succeed in satisfying entertainment goal)
3. Unsuccessful search for connection between 'plan success' postdiction and 'depression' affect
4. Successful search for connection between alternate 'plan failure' postdiction and 'depression' affect

The two crucial things to note are: (1) precisely the same content inferences are made in the same circumstances by both Recencies and Perseverers; the only difference is when they make them. It is only that difference that leads to the differences in eventual interpretation. (2) Perseverer behavior on text 1F and Recency behavior on text 1B are almost identical to each other; reciprocally, Perseverer behavior on text 1B and Recency behavior on text 1F are almost identical.

### 3.0 Operation of the STRATEGIST prototype

The following represents actual annotated run-time output of the STRATEGIST program. First we examine STRATEGIST's behavior as a Perseverer. The input to the program is the Conceptual Dependency representation (Schank & Abelson, 1977) of the following story:

[2] Melissa began to cry. Tyler had just asked her to marry him.

:processing as perseverer

:the story is:

(EXPEL (ACTOR MELISSA) (OBJECT TEARS))  
(MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)  
(TO MELISSA) (FROM TYLER))

:processing next concept:

(EXPEL (ACTOR MELISSA) (OBJECT TEARS))

:attempting inference generation

:inferring from: MELISSA  
:no context found  
:default inference selected:  
(A-HAPPINESS (PLANNER MELISSA))  
(DO-X (ACTOR MELISSA))  
(MENT-ST (ACTOR MELISSA) (VALUE POS))

:inferring from: TEARS  
:no context found  
:default inference selected:  
(DO-X (ACTOR ?ACTOR0))  
(MENT-ST (ACTOR MELISSA) (VALUE NEG))  
(EXPEL (ACTOR MELISSA) (OBJECT TEARS))

STRATEGIST searches for existing inferences (i.e. the context which constrains the current inference generation) which might connect with the inferences to be generated. No existing, applicable context is found, so STRATEGIST searches for any inferences associated with TEARS in the context of an EXPEL. It finds a default inference that someone has done something which made Melissa unhappy and caused her to cry tears of sadness. (Note that if this inference later turns out to be incompatible with some subsequent inference, it may be supplanted [Granger, 1980] and an alternative inference used.)

:end of inference generation

:processing next concept:

(MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)  
(TO MELISSA) (FROM TYLER))

All inferencing has been completed for the first conceptualization, so STRATEGIST begins inference generation for the next conceptualization. The inferences generated from the first conceptualization provide the context in which the next conceptualization will be interpreted.

:attempting inference generation

:inferring from: TYLER  
:no context found  
:default inference selected:  
(A-HAPPINESS (PLANNER TYLER))  
(DO-X (ACTOR TYLER))  
(MENT-ST (ACTOR TYLER) (VALUE POS))  
  
:inferring from: PROPOSE-MARRIAGE  
:context found  
:possible inferences are:  
(MTRANS (ACTOR TYLER) (MOBJ  
PROPOSE-MARRIAGE)  
(FROM TYLER) (TO MELISSA))  
(MTRANS (ACTOR MELISSA) (MOBJ ACCEPT)  
(FROM MELISSA) (TO TYLER))  
(MENT-ST (ACTOR TYLER) (VALUE POS))  
(MENT-ST (ACTOR MELISSA) (VALUE POS))

(MTRANS (ACTOR TYLER) (MOBJ  
PROPOSE-MARRIAGE)  
(FROM TYLER) (TO MELISSA))  
(MTRANS (ACTOR MELISSA) (MOBJ REJECT)  
(FROM MELISSA) (TO TYLER))  
(MENT-ST (ACTOR TYLER) (VALUE NEG))  
(MENT-ST (ACTOR MELISSA) (VALUE NEG))

STRATEGIST finds that an applicable context does exist for PROPOSE-MARRIAGE, so it looks at the possible predictive inferences for PROPOSE-MARRIAGE: that Tyler proposes, Melissa accepts, and both are happy, or that Tyler proposes, Melissa rejects his offer, and both are unhappy.

:found matching inference  
:resulting merged inference is:  
(MTRANS (ACTOR TYLER) (MOBJ  
PROPOSE-MARRIAGE)  
(FROM TYLER) (TO MELISSA))  
(MTRANS (ACTOR MELISSA) (MOBJ REJECT)  
(FROM MELISSA) (TO TYLER))  
(MENT-ST (ACTOR TYLER) (VALUE NEG))  
(MENT-ST (ACTOR MELISSA) (VALUE NEG))  
(EXPEL (ACTOR MELISSA) (OBJECT TEARS))

The inference that Melissa rejected Tyler's proposal and she is unhappy connects with the previously made predictive inference that Melissa was crying because someone did something which made her unhappy. The new inference chain which results is then stored as a predictive inference which will be applied to future inference generation.

:end of inference generation

:end of processing

:final representation is:

(A-HAPPINESS (PLANNER MELISSA))  
(DO-X (ACTOR MELISSA))  
(MENT-ST (ACTOR MELISSA) (VALUE POS))  
  
(A-HAPPINESS (PLANNER TYLER))

```
(DO-X (ACTOR TYLER))
(MENT-ST (ACTOR TYLER) (VALUE POS))
```

```
(MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)
  (FROM TYLER) (TO MELISSA))
(MTRANS (ACTOR MELISSA) (MOBJ REJECT)
  (FROM MELISSA) (TO TYLER))
(MENT-ST (ACTOR TYLER) (VALUE NEG))
(MENT-ST (ACTOR MELISSA) (VALUE NEG))
(EXPEL (ACTOR MELISSA) (OBJECT TEARS))
```

STRATEGIST finishes processing the story and prints the inferences it has made. The first two inference chains listed above show that both Tyler and Melissa had goals of achieving happiness, but their goals were unfulfilled. The last chain indicates the order of actual events as STRATEGIST inferred them: that Tyler proposed, Melissa said "no", both were unhappy, and Melissa cried.

STRATEGIST can be told to apply the Recency strategy by changing the value of a parameter. This in turn invokes a process which postpones the processing of input until the end of the input is detected. We now examine the operation of STRATEGIST as a Recency:

```
:processing as recency

:the story is:
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))
  (MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)
    (TO MELISSA) (FROM TYLER))

:processing next concept:
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))

:leaving loose end

:processing next concept:
  (MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)
    (TO MELISSA) (FROM TYLER))

:leaving loose end

:end of input story
```

Behaving as a Recency, STRATEGIST postpones high-level inference generation until no input remains to be processed. It then begins to generate inferences from the input conceptualizations in last-in-first-out order.

```
:processing previous loose end:
  (MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)
    (TO MELISSA) (FROM TYLER))

:attempting inference generation

:inferring from: TYLER
:no context found
:default inference selected:
  (A-HAPPINESS (PLANNER TYLER))
  (DO-X (ACTOR TYLER))
```

```
(MENT-ST (ACTOR TYLER) (VALUE POS))
```

The default inference for TYLER becomes the context for further inferencing.

```
:inferring from: PROPOSE-MARRIAGE
:context found
:possible inferences are:
  (MTRANS (ACTOR TYLER) (MOBJ
    PROPOSE-MARRIAGE)
    (FROM TYLER) (TO MELISSA))
  (MTRANS (ACTOR MELISSA) (MOBJ ACCEPT)
    (FROM MELISSA) (TO TYLER))
  (MENT-ST (ACTOR TYLER) (VALUE POS))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))

  (MTRANS (ACTOR TYLER) (MOBJ
    PROPOSE-MARRIAGE)
    (FROM TYLER) (TO MELISSA))
  (MTRANS (ACTOR MELISSA) (MOBJ REJECT)
    (FROM MELISSA) (TO TYLER))
  (MENT-ST (ACTOR TYLER) (VALUE NEG))
  (MENT-ST (ACTOR MELISSA) (VALUE NEG))

:found matching inference
:resulting merged inference is:
  (A-HAPPINESS (PLANNER TYLER))
  (MTRANS (ACTOR TYLER) (MOBJ
    PROPOSE-MARRIAGE)
    (FROM TYLER) (TO MELISSA))
  (MTRANS (ACTOR MELISSA) (MOBJ ACCEPT)
    (FROM MELISSA) (TO TYLER))
  (MENT-ST (ACTOR TYLER) (VALUE POS))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))
```

STRATEGIST now finds the same two possible inferences for PROPOSE-MARRIAGE that it did while behaving as a Perseverer, but this time finds that the previously made explanatory inference that Tyler wanted to be happy connects with the possible inference that Melissa accepts Tyler's proposal and both actors are happy. The resulting connected inference chain is stored in memory and serves as the context for later inferencing. Contrast this with STRATEGIST's behavior as a Perseverer, in which it was inferred that Melissa rejected Tyler's proposal and was unhappy because of the existing context that Melissa was crying because of her unhappiness.

```
:end of inference generation

:processing previous loose end:
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))

:attempting inference generation

:inferring from: MELISSA
:context found
:possible inferences are:
  (A-HAPPINESS (PLANNER MELISSA))
  (DO-X (ACTOR MELISSA))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))

:found matching inference
```

```

:resulting merged inference is:
  (A-HAPPINESS (PLANNER TYLER))
  (MTRANS (ACTOR TYLER) (MOBJ
PROPOSE-MARRIAGE)
    (FROM TYLER) (TO MELISSA))
  (A-HAPPINESS (PLANNER MELISSA))
  (MTRANS (ACTOR MELISSA) (MOBJ ACCEPT)
    (FROM MELISSA) (TO TYLER))
  (MENT-ST (ACTOR TYLER) (VALUE POS))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))

STRATEGIST finds that the default goal of
Melissa's wanting to be happy coincides with
the existing context.

```

```

:inferring from: TEARS
:context found
:possible inferences are:
  (DO-X (ACTOR ?ACTOR0))
  (MENT-ST (ACTOR MELISSA) (VALUE NEG))
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))

  (DO-X (ACTOR ?ACTOR0))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))

```

```

:found matching inference
:resulting merged inference is:
  (A-HAPPINESS (PLANNER TYLER))
  (MTRANS (ACTOR TYLER) (MOBJ
PROPOSE-MARRIAGE)
    (FROM TYLER) (TO MELISSA))
  (A-HAPPINESS (PLANNER MELISSA))
  (MTRANS (ACTOR MELISSA) (MOBJ ACCEPT)
    (FROM MELISSA) (TO TYLER))
  (MENT-ST (ACTOR TYLER) (VALUE POS))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))

```

Here, STRATEGIST finds that the possible inference of Melissa's crying tears of joy coincides with the existing context. Once again, contrast this with the Perseverer behavior above in which STRATEGIST selected the default inference of Melissa's crying tears of sadness.

```

:end of inference generation

:end of processing
:final representation is:

  (A-HAPPINESS (PLANNER TYLER))
  (MTRANS (ACTOR TYLER) (MOBJ PROPOSE-MARRIAGE)
    (FROM TYLER) (TO MELISSA))
  (A-HAPPINESS (PLANNER MELISSA))
  (MTRANS (ACTOR MELISSA) (MOBJ ACCEPT)
    (FROM MELISSA) (TO TYLER))
  (MENT-ST (ACTOR TYLER) (VALUE POS))
  (MENT-ST (ACTOR MELISSA) (VALUE POS))
  (EXPEL (ACTOR MELISSA) (OBJECT TEARS))

```

STRATEGIST ends processing, its final representation indicating that all inferences were connected. The inferred explanation for the events related to STRATEGIST is that both

Tyler and Melissa wanted to be happy, Tyler asked Melissa to marry him, Melissa said "yes", both actors were happy, and Melissa cried.

#### 4.0 Interesting Observations

One problem that this model addresses is determining which of two equally plausible and parsimonious interpretations will be selected. Along these lines, Schank, et. al. explain such misunderstanding in verbal communication by "...maintaining that deriving a point is a part of processing, specifically related to the choice of an 'inference path.' Understanders choose to process idiosyncratically" (Schank, et. al., 1982, p. 263). This explanation of deriving a point agrees with our theory of inference paths. However, rather than believing understanders' processing to be idiosyncratic, our model predicts that individuals will tend to follow a single strategy consistently, rather than arbitrarily switching from path to path.

Behavior of Recencies and Perseverers is notable because it supports the theory that the strategies use the same component processes to interpret text. For example, both strategies see a single interpretation of the text. In some cases, when the alternate interpretation is pointed out, subjects will protest that the alternate interpretation is implausible, based on the way events were presented, regardless of the strategy they employed. Our experiments also indicate that readers using the two different strategies will reverse their interpretation of events if the order of events in the text is reversed.

Readers using either strategy can be forced to switch to the opposite strategy. For example, the typical experimental method for studying inference decisions presents a text to the subject a sentence at a time, and asks the subject what inferences were made after each sentence. If a Recency is given text one line at a time, so that no cues about the existence of further events can be used, his interpretation will be the same as a Perseverer's, even for those stories which would normally result in a different interpretation. Thus, the data collected will not reveal the different strategies. It is only when subjects are allowed to read a full text, and not forced to make inferences by the experimenter, that the different strategies can be observed. In fact, previous researchers (e.g. Rumelhart, 1981; Seifert, Robertson, & Black, 1982) have used a line-at-a-time methodology for studying when and which inferences will be made. In Rumelhart's (1981) experiments, subjects read a text a sentence at a time. After each sentence, the subject was asked for his current interpretation of the text. When Rumelhart compared the interpretations of the texts by subjects who

read the text a sentence at a time to the interpretations of subjects who read the full text and were not asked for their interpretations until after completing the text, he found that the subjects who read the texts all at once "showed somewhat more variability in their interpretations", which he attributed to "more careless reading on the part of the subjects offering an interpretation only at the end" (Rumelhart, 1981).

What happens when one's usual strategy cannot be used? It is possible that the Inference Manager has several sets of rules from which to choose, and that other sets of rules are invoked when the "default" set fails. For example, a Perseverer who doubted his initial interpretation would use the Recency strategy to discover a new interpretation. Our experimental evidence regarding new interpretations in response to requestioning (Granger & Holbrook, 1983) leads us to reject this hypothesis. Instead, we theorize that an individual's Inference Manager has only one set of rules, certainly more complicated than those which we have described, with many "if/then/else" alternatives.

One might suspect that the only difference between the two strategies is which inference is chosen as the default inference. The evidence does not support such a theory; if a Recency were making the original default inference after the first concept was presented, but also making default inferences for later concepts and simply choosing the later concepts when defaults conflict, then Recencies would presumably have little trouble recognizing the Perseverer's interpretation as an alternative interpretation. As discussed earlier, this is not the case, nor have reaction-time tests on false recognition items suggested otherwise (Granger & Holbrook, 1983).

Perseverers and Recencies are only two points in a range of strategies. An extreme Perseverer makes inferences based on a preconceived context. This strategy is a kind of paranoid understanding. An extreme Recency will not make inferences, and will not be able to understand text which requires any higher-level inferences. Still other readers exhibit behavior akin to both Recency and Perseverer behavior. We call these readers Deferrrers; at present, little is understood about the strategies used by Deferrrers.

## 5.0 Summary and Conclusions

We have presented evidence for processes of story comprehension which include the set of rules used by most story understanding programs, and an additional set of rules which accounts for interpretations which these programs would not be able to make.

Our prototype model is far from finished. One limitation is that it uses simplistic representations. For example, the representations do not include knowledge about which plans are appropriate for a goal, nor do they include knowledge about the possible conditional outcomes of plans. Questions re-presented to STRATEGIST will not result in another interpretation of events. These are all extensions to the system which are planned for the future.

STRATEGIST is primarily a model of human understanding. There are still many questions to be answered about how people interpret text. Our experiments have yet to reveal all of the different strategies used by readers. We have studied evaluation metrics and processes, as well as some of the rules which apply the processes. Future work will focus on specifying more rules, and more carefully defining and ordering those rules which we have described here. We also hope to study the application of these strategies applied to longer texts of many different genres. This work will not only involve observation of human subjects; the extended STRATEGIST will be a test-bed which will allow us to study inference processes and new rules which apply those processes.



## REFERENCES

- Graesser, A.C. A question answering method of exploring prose comprehension: an overview. Proceedings of the Third Annual Conference of the Cognitive Science Society, Berkeley, California, 1981.
- Granger, R.H. When Expectation Fails: Towards a self-correcting inference system. Proceedings of the First National Conference on Artificial Intelligence, Stanford, California, 1980.
- Granger, R.H. Directing and Re-Directing Inference Pursuit: Extra-Textual Influences on Text Interpretation. Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, British Columbia, 1981.
- Granger, R.H., & Holbrook, J.K. Perseverers, Recencies, and Deferrers: new experimental evidence for multiple inference strategies in understanding. Proceedings of the Fifth Annual Conference of the Cognitive Science Society, Rochester, New York, 1983.
- Rumelhart, D.E. Understanding understanding (Report No. 100). San Diego, California: Center for Human Information Processing, University of California at San Diego, 1981.
- Schank, R.C. & Abelson, R.P. Scripts, Plans, Goals, and Understanding. Hillsdale, New Jersey: Erlbaum, 1977.
- Schank, R.C., Collins, G.C., Davis, E., Johnson, P.N., Lytinen, S., & Reiser, B.J. What's the Point? Cognitive Science, 1982, 6, 255-275.
- Seifert, C.M., Robertson, S.P. & Black, J.B. On-Line processing of pragmatic inferences (Report No. 15). New Haven, Connecticut: Cognitive Science Program, Yale University, 1982.
- Wilensky, R. Understanding goal-based stories (Report No. 140). New Haven, Connecticut: Department of Computer Science, Yale University, 1978.
- Wilensky, R. Planning and understanding. Reading, Massachusetts: Addison-Wesley, 1983.