DIRECTING AND RE-DIRECTING INFERENCE PURSUIT: Extra-textual influences on text interpretation

Richard H. Granger, Jr.

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

ABSTRACT

Understanding a text depends on s reader's ability to construct a coherent interpretation that accounts for the statements in the text. However, a given text does not always imply a unique coherent interpretation. In particular, readers can be steered away from an otherwise plausible explanation for a story by such extra-textual factors as the source of the text, the reading purpose, interruptions during reading, or repeated re-questioning of the reader. Some of these effects have been observed in experiments in cognitive psychology (e.g., Black [1980]). This paper presents a computer program called MACARTHUR that can vary both the depth and direction of its inference pursuit in response to re-questioning, resulting in a series of markedly different interpretations of the same text*

1) INTRODUCTION

Consider the following story:

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

Why did the ambassador go to Albania? People in informal experiments most often answer that he may have simply gone there as part of his vacation. However, when the same question is repeated, they generate alternative explanations, such as the following:

1. There could have been some secret political meeting there.

2. There might have been plane trouble; say, an emergency landing to fix a fuel leak.

3. Maybe he juat wanted to avoid reporters on his vacation.

This research was supported in part by the Naval Ocean Systems Center under contract N66001-80-C-0377.

The text presents an explanation on the surface (that the ambassador was on vacation), which is adequate to serve as an interpretation of the events in the story. Yet re-probing with the same question has the effect of causing people to generate new and "slanted" interpretations, reading more into the passage than before. In a related series of informal experiments, people were told different "sources" of the text; in particular, they were either told that it was excerpted from the New York Times, an Agatha Christie novel, Cosmopolitan magazine, a grammar-school history textbook or a Jimmy Stewart movie. Their interpretations of the text varied significantly depending on the stated text source.

These observations about people's reading behavior agree with experiments in cognitive psychology in which varying the stated reading purpose (e.g., Black [1980], Frederiksen [1975]), and interposing questions about the text (e.g., Rothkopf and Bisbicos (1967], Anderson and Biddle [1975]) resulted in differences in the inferences that were made by the readers, as evidenced by tests for false recognition of statements corresponding to inferences from the text.

This paper presents a program called MACARTHUR which is able to redirect its own inference processes when a question about a text is re-asked repeatedly. MACARTHUR demonstrates its successive interpretations by generating English answers to questions about the text. For example, after reading a version of the above story [1], MACARTHUR responds in English to the following sequence of questions:

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

Q) Are you sure? 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.

Most existing natural language systems (e.g., Cullingford [1978], Wilensky [1978], Charniak [1978]) do not account for people's ability to make different inferences depending on external factors such as re-probing. MACARTHUR models this observed human ability by using a hierarchy of understanding processes, some of which "steer" the direction of pursuit of others.

2. BACKGROUND: CORRECTING ERRONEOUS INFERENCES

2.1 Maintaining a connected representation

MACARTHUR is an extension of the ARTHUR system, described in Granger (1980s, 1980b]. That paper pointed out that the process of mapping a story onto a representation is not always straightforward, but rather may require the generation of a number of intermediate representations which are supplanted by the time the final story representation is complete. For example, consider the following simple story, taken from Granger [1980b]:

[2] Geoffrey Huggins walked into the Roger Sherman movie theater. Re went up to the balcony, where Willy North was waiting with a gram of cocaine. Geoff paid Willy in large bills and left quickly.

Why did Geoff go into the movie theater? Most people infer that he did so in order to buy some coke, since that was the outcome of the story. The alternative possibility, that Geoff went to the theater to see a movie and then coincidentally ran into Willy and decided to buy some coke from him, seems to go virtually unnoticed. On the basis of pure logic, either of these inferences about Geoff's intentions is equally plausible. However, people overwhelmingly infer the former explanation, i.e., that he went to the theater intending to buy coke.

The problem is that the most plausible initial inference from the story's first sentence is that Geoff did go inside to see a movie. Hence, selection of the correct inference requires rejection of this initial inference. The point is that Geoff did not change his mind about why he went into the theater; he had a single intention from the outset. Rather it is we, the readers, who must change our minds about our initial inference of Geoff's intention. ARTHUR is able to understand stories like [2] because of its ability to re-evaluate and supplant its own initial inferences in light of subsequent information in a story.

Story [1] presents us with a different but related difficulty. There is no strong inference in [1] that causes readers to supplant their initial inference; the "vacation" explanation adequately accounts for the statements in the story. However, readers can be steered away from this explanation by external factors such as varying the text source or repeating the same question. The problem MACARTHUR addresses is how external factors like re-probing can affect the inferences a reader produces.

Story [1] describes two events: the ambassador leaves for Pakistan on vacation, and then he makes an unscheduled stop in Albania. MACARTHUR's interpretation of the first event simply contains both the stated action of physically leaving and the stated goal of going on a vacation. Since changing your location can be part of a known plan for vacationing, MACARTHUR assumes that that is the connection between them.

MACARTHUR then attempts to find an explanation for the next event, stopping in Albania. It first checks whether the event can be interpreted as part of a plan in service of the already existing goal of vacationing. Since going to Albania is also a change of location, MACARTHUR assumes that this action too can be interpreted as being part of the vacation. In answer to a question about why the ambassador went to Albania, then, MACARTHUR uses this explanation it has constructed to generate the answer that the trip was part of the vacation.

2.2 Re-directing inference pursuit

MACARTHUR has arrived at a connected representation for the story, and has answered a question about it. Now the program is "re-probed" with the same question, causing it to re-evalutate its initial explanation for the ambassador's going to Albania. MACARTHUR finds that Albania, in addition to being a location (and therefore the possible site of a vacation), is also a political entity. Hence, MACARTHUR infers that a trip to Albania could be part of a plan to meet with members of Albania's government.

Now MACARTHUR has a new plausible explanation for the ambassador's action of going to Albania, but his action of going to Pakistan is still explained only by the original "vacation" inference. This is • an example of a disconnected representation, i.e., one in which the explanations for the events of a story are not connected to each other. The current explanation would result in MACARTHUR answering that the ambassador went to Albania to confer with the government, but that he went to Pakistan to vacation. Although there is nothing wrong with this on the basis of pure logic, it does not at all correspond to the inferences people make at this point.

It is important to note that the "vacation" explanation is not explicitly contradicted by the new explanation; it just fails to be connected to it. Granger [1980b] hypothesized that people attempt to connect the pieces of a representation according to the "parsimony principle", which states that the best goal inference is one which accounts for the most actions of an actor. Hence, the best story representation is one which contains the fewest number of context inferences to explain the most events in the story.

The correct representation should, therefore, connect the ambassador's intention of conferring with the Albanian government together with both his action of travelling to Albania, and his announced vacation trip to Pakistan. To achieve this connection, MACARTHUR infers that the ambassador may have had a constraint on his goal, of keeping his Albanian meeting secret. A known method of keeping an event secret is to create a "cover story" for the event, i.e., an alternative explanation that can account for the covert action. In this case, MACARTHUR assumes that the ambassador announced that he was going on vacation as a cover story for his trip to Albania. (The structure of cover stories and MACARTHUR's ability to recognise them is elaborated in Section 3.2.)

MACARTHUR's representation of the story now consists of a single inference about the ambassador's intentions (he wanted to confer with the Albanian government in secret), and a plan in service of that goal (going to Albania); along with the announced "cover story" about going on vacation, in service of the "secrecy" constraint. The initial goal inference that the ambassador actually intended to go on vacation has now been supplanted: it is no longer considered to be the explanation for the events in the story.

2.3 Experimental Evidence of Effects *on Inference*

A number of experiments in cognitive psychology have demonstrated that people's reading behavior can be significantly affected by conditions independent of the content of the text itself (e.g., Anderson and Biddle [1975], Frederiksen [1975], Black [1980], Haberlandt and Bingham [1978], Rothkopf and Bisbicos [1967]). For instance, reading an assigned text on which you may be tested can result in quite different behavior than reading the same text out of pure interest in the topic, reading it to attempt to criticise it, reading it becauae your brother wrote it, etc.

In particular, Frederiksen [1975] and Black [1980] have supported the hypothesis that people make more inferences, and different types of inferences, depending on the purpose they are given by the experimenter for reading the text. Black for example notes that people make more inferences when asked to read a passage in preparation for a subsequent memory test than when they are asked to read it in order to simply rate it for "comprehensibility". Since the latter task is "shallower", it requires less of what Black terms "elaborative" inference (inferences above and beyond strictly logical deductions).

Black also notes that when subjects were asked to read the passage for the purpose of later writing an essay that uses the main point of the passage, they tended to make more inferences that related specifically to the main point of the passage, although they did not make more elaborative inferences overall than the subjects who were to be tested for their memory for the passage. This strongly implies that a stated purpose for reading a text can cause a reader not only to suppress or accelerate his inference processes, but also to steer the direction of pursuit of the inferences that are generated.

To illustrate this another way, consider using story [1] (the ambassador story) as the text for a similar experiment. Since this text is more difficult to arrive at a single explanation for, the results might show "shallow" readers inferring the "naive" interpretation of the story, i.e., that the ambassador was just going on vacation, while "deeper" readers might infer one of the alternative interpretations, e.g., that there was a secret meeting in Albania. This would demonstrate that deeper reading tasks could result in different inferences, not just more inferences, being generated from a single text.

MACARTHUR's ability to re-direct its own inferences is designed to model people's observed reading behavior in similar tasks. The model is intended to provide a test-bed for comparing implementations of our theories about people's reading behavior with actual experimental evidence. Towards this end, section 5 of this paper proposes some possible new experiments and possible extensions to MACARTHUR.

3. OPERATION OF THE MACARTHUR PROGRAM

3.1 Annotated run-time output

The following represents actual annotated run-time output of the MACARTHUR program. The input to the program is the Conceptual Dependency representation (Schank and Abelson [1977]) of the following story:

[3] 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 (extremely abridged) 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 R00M1) (FROM ROOM0))
 PC1: (D-PROX (PLANNER FITZ)
              (OBJECT FITZ) (TO R00M1))
 PATH1: ($CO-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 getting 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 R00M1))

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].) (Note: The "explanation shape" reported here as "pursue-desired-state" is one of four such categories of explanations that MACARTHUR knows about. These are explained in Section 3.2.)

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 ROOMI))

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'8 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.

3.2 Note: The shape of things to come

When MACARTHUR attempts to re-explain a story because of re-probing, it uses a classification scheme for discerning the "shape" of an explanation, which aids it in selecting alternative explanations. This scheme is still the subject of ongoing research, but it has proved useful in MACARTHUR's explanation selection algorithm. Following is a list of some of the shapes MACARTHUR currently knows about. This is not intended to be a complete list, it simply reflects the present state of our analysis:

1. Pursue-desired-state: This refers to simple goal pursuit, i.e. a story in which a character has a goal and performs plans in service of that goal.

2. Avoid-undesired-state: A character may not have a specific goal or desired state, but rather is acting out plans that are in service of the avoidance of a particular undesired state, such as sleepiness (for which a remedy is to ingest coffee or other stimulants), hunger (remedies include doing something distracting like reading, or taking diet pills, or even going to sleep), etc.

3. Accident-reaction: A character may be involved in some events that unintentionally hinder his goals. The character's subsequent actions may include attempts to investigate the cause of the accident; overcoming the accident by re-planning and re-acting; abandoning or postponing the goal; or simply trying again.

4. Cover-stories: A character may have a goal that he wishes to achieve secretly. If he cannot simply avoid being observed, then he may

construct a "cover story"; i.e., an alternative connected explanation for his actions vhich can serve as an "alibi" to any observers. Complete understanding of such stories involves the ability to maintain separate belief spaces for different characters, and to recognise deception via conflicting beliefs held by different characters.

Now we will illustrate how these explanation shapes can give rise to a series of alternative interpretations of stories. Recall story [1]:

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

The four alternative explanations given for this story at the beginning of this paper can now be categorised by explanation shape:

1. He may have gone there as part of his vacation. (PURSUE-DESIRED-STATE)

2. There could have been some secret political meeting there. (COVER-STORY)

3. There might have been plane trouble; say an emergency landing. (ACCIDENT-REACTION)

4. Maybe he just wanted to avoid reporters on his vacation. (AVOID-UNDESIRABLE-STATE)

Now recall story [3]:

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

Following are four differently-shaped explanations for this story, (two of which correspond to explanations generated by MACARTHUR in section 3):

1. Maybe he wanted to make some warm milk to help him get to sleep. (PURSUE-DESIRED-STATE)

2. Maybe he wanted to make some coffee to help him stay awake. (AVOID-UNDESIRABLE-STATE)

3. Maybe he heard something fall down in there and he went to investigate. (ACCIDENT-REACTION)

4. Maybe he actually had some secret reaaon for going in there, so he yawned to pretend he was tired. (COVER-STORY)

The next section describes the processes by which MACARTHUR uses its knowledge of explanation shapes to generate alternative interpretations of a story when re-probed.

4. THE PROCESSES UNDERLYING INFERENCE RE-DIRECTION

4.1 Hypothesis-selection and hypothesis-pursuit

Understanding a story or situation often requires us to generate a hypothesis about the goals of the participants, on the basis of their observed actions (selection); and then to attempt to continue to explain subsequent actions in terms of the hypothesized goals (pursuit). Any sufficiently difficult text can suggest multiple alternative explanations, and the reader must select one and pursue it, but he must also be ready to supplant an initial hypothesis with a new one in light of subsequent information.

4.2 Pursuing hypotheses

MACARTHUR's understanding algorithm essentially consists of attempting to explain all of the conceptualisations in a text in terms of a minimum number of hypotheses, in accordance with the "parsimony principle" (Granger [1980b]). For example, in attempting to connect up a new story event with an existing hypothesis about a character's goal, MACARTHUR begins by pursuing the goal hypothesis, attempting to integrate the event into the hypothesis, via the following steps:

1. Search for existing "indexed inference paths" (see Granger [1980a]), which correspond to pre-stored inferential paths that can serve to connect the hypothesis with the statement;

2. Search known preconditions of the goal for matches with the event;

3. Generate bottom-up indexed inferences from the event to see if any can match existing inference paths to the goal.

4.3 Supplanting hypotheses

If these steps fail, then MACARTHUR may decide to supplant the current hypothesis with a new one; i.e., abandon its pursuit of the hypothesis, and select an alternative hypothesis to explain the "recalcitrant" event. This process is described in detail in Granger [1980b].

1. Generate alternative hypothesis by choosing different explanation-shape;

2. Attempt to explain previous events in terms of new hypothesis, via rules of pursuit.

4*4 Leaving loose ends

If the procedures above fail to re-explain the events in terms of the new hypothesis, then MACARTHUR is in the situation of having two competing hypotheses, the initial one and the new one, neither of which can explain all of the events in the story. This implies that each of the hypotheses has run into one or more "recalcitrant" events. In this caae, MACARTHUR tentatively admits defeat by reverting back to the initial hypothesis, and marking its recalcitrant event as a "loose end" relative to the hypothesis; i.e., an event that

cannot be explained by the existing hypothesis. Similarly, the alternative hypothesis is maintained, and its recalcitrant events are also marked as loose ends relative to this hypothesis:

1. Mark recalcitrant events as loose ends relative to new hypothesis;

2. Reinstate initial hypothesis;

3. Mark its recalcitrant events as loose ends relative to it.

When MACARTHUR leaves a loose end, it then has a disconnected explanation. Hence, loose ends result in MACARTHUR's interpretation of the story being less than maximally parsimonious. (This paper has not shown any examples of MACARTHUR leaving a loose end.)

4.5 Re-shaping explanations

Once MACARTHUR has arrived at a complete explanation for a story, then it can answer questions that refer to the explanation, such as "why"-questions asking about characters goals. MACARTHUR's processes of searching its explanation to provide an answer, and then generating that answer in English, are similar to those used by ARTHUR, as described in Granger [1980a]. However, MACARTHUR has the additional ability to re-explain a story in response to repetition of a question. This requires selection of a new hypothesis as a starting point for the new explanation. MACARTHUR does this by first selecting a shape different from the current explanation shape, and then letting that shape suggest a new goal hypothesis:

1. Choose alternative explantion shape;

2. Use new shape to generate alternative goal hypothesis;

3. Attempt to replace existing hypothesis with new hypothesis via rules of supplanting.

Table 1 (below) illustrates a few examples of the knowledge MACARTHUR uses to construct a goal explanation for a given event conforming to particular shapes. For each of these event/shape pairs, examples abound. For example, the pair <PTRAMS/COVER-STORY> corresponds to the "secret meeting" explanation of story [1], the pair <PROPEL/ACCIDENT-REACTION> could arise if Z PROPELed X towards W, but it went to Y instead; 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/ACCIDENT-REACTION> could arise if Z INGESTed something that he wouldn't have INGESTed had he known what it was; e.g., 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.

	PURSUE-STATE	AVOID-STATE	COVER-STORY	ACCIDENT-REACTION
(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 1: Shaping the Interpretation of Events

5. CONCLUSIONS: PROPOSALS FOR FURTHER RESEARCH

5.1 focperjmeiUs, on inference re-direction

Black's [1980] experiments on the effects of reading purpose on memory for text assumed that the task of rating the comprehensibility of a text was "a 'shallow' task", preparing for a memory test was "a 'deeper' task", and preparing for an essay test in which the subjects would have to make use of the main point of the text was "a 'deepest' task" [p. 20]. Black's initial prediction was basically that the "deeper" the reading purpose, the greater the number of inferences the subject would produce, as evidenced by the number of false recognitions exhibited on tested inference items.

The actual results of the experiment indicated that the memory task caused the most false recognitions of inference items, while the essay task came second and the comprehensibility task came lowest, as expected. A post-hoc analysis of the recognition test items revealed that the essay task caused significantly more false recognitions than the other two groups on inference items which were "related- to the main point'' of the JJ&£X, even though the number of false recognitions overall (i.e., including items both related and unrelated to the main point) was lower for the essay task than for the memory task.

In other words, the experiment was looking for a monotoniesIly increasing effect of more inferences corresponding to "deeper*' processing. However, what it found was a difference in not only the "depth", but also in the "direction" of inferences generated. In particular, Black acknowledges the existence of "main-point oriented" processing in the essay task which did not appear in the other two tasks.

Consider a similar set of experiments performed using more difficult stories, i.e., stories that are less strongly connected to a single main point than the essays used in Black's study. For example, texts like [1] and [3] in this paper could be used. As noted by many researchers (e.g. Haberlandt and Bingham [1978], Black and Bern [1980]), readers tend to work at finding connections among sentences in a text, even when such connections are not obvious. Hence we predict that subjects would dutifully generate connective inferences to explain the sentences in a non-straightforward text like [1] and [3], but since there are a number of different alternative interpretations of texts like these, different explanations might be produced by different subjects, perhaps as a function of different types of external factors such as reading purpose, text source, interposed questions and re-probing. For example, in a reading-purpose experiment the "shallower" readers might generate a "naive" interpretation of a difficult text; while deeper readers might generate not just more inferences but different inferences, corresponding to their significantly different interpretation of the text.

We respectfully propose such a set of experiments, designed around non-straightforward texts, and making use of other types of extra-textual factors than just reading purpose; in particular, the effects of interposed questions and re-probing. Some questions that might be resolved by these experiments include the following:

1. do peoples' alternative explanations correspond well with the classification of explanations into different "shapes" (as proposed in section 3.2), and/or with other inference classifications such as the three binary classifications (derived/elaborative, local/global and few-steps/many-steps) proposed by Black [1980]?

2. do certain categories of explanations correspond to longer reading times?

3. do certain categories of explanations correspond to better recognition or recall for the stories later on?

4. are there any observed regularities in the order in which people generate certain categories of explanations in response to re-probing?

There are certainly many other interesting issues dealing with people's inferencing and story-explanation abilities; this list is just meant to be suggestive of some issues that might be able to be resolved by the experiments proposed here.

5.2 Proposed extends to MACARTHUR

MACARTHUR can generate alternative explanations for a given text in response to repeated questioning about a particular point in the text. This paper has mentioned a number of types of external factors that can influence the inferences generated from a text, and thereby can give rise to alternative explanations for a text, e.g., text source, interposed questions, re-probing, and varying reading purpose.

It is not clear how to construct a computer model corresponding to the situations of giving a reader a specified reading purpose nor a specified text source, since most of the task would be to build enough knowledge into the program to model people's extensive knowledge of different particular text sources and of the intricacies associated with certain reading purposes. However, the issue of interposing questions at various points in the middle of the text is one which potentially could be modelled. We plan to investigate the literature of experiments in this area, and hope to extend MACARTHUR to model the effects of interposed questions on people's inference generation and explanation of stories.

ACKNOWLEDGEMENT

Most of the ideas in this paper benefited from lengthy discussions with Jim Meehan, who also made valuable suggestions for the various drafts.

REFERENCES

[1] Anderson, R.C. and Biddle, W.B. (1975) On asking people questions about what they are reading. In G. H. Bower (Ed.) The psychology of learning and motivation. Vol 9. Academic Press, New York.

[2] Black, J. (1980). The Effects of Reading Purpose on Memory for Text, Cognitive Science TR #7, Yale University, New Haven, Conn.

[3] Black, J. and Bern.H. (1980). Causal Coherence and Memory for Events in. Narratives, Cognitive Science TR #3, Yale University, New Haven, Conn.

[4] Bower, G.H., Black, J.B. and Turner,T. (1979). Scripts in memory for text. Cognitive Psychology, ii, 177-220.

[5] Charniak, E. (1978). On the use of framed knowledge in language comprehension. Computer Science TR #137, Yale University, New Haven, Conn.

[6] Cullingford, R. (1978). Scrjpt Application: Compter Understanding of Newspaper Stories Ph.D. Thesis, Yale University, New Haven, Conn.

[7] DeJong, G.F. (1979). Skimming Stories in Real Time: An Experiment in Integrated Understanding- Ph.D. thesis, Dept. of Computer Science, Yale University, New Haven, Conn.

[8] Frederiksen, C.H. (1975). Effects of context-induced processing operations on semantic information acquired from discourse. Cognitive Psychology. 2, p. 371-458.

[9] Granger, R. (1980a). Adaptive Understanding: Correcting Erroneous Inferences, Ph.D. Thesis, Yale University, New Haven, Conn.

[10] Granger, R. (1980b). When Expectation Fails: Towards a Self-correcting inference system, Proceedings of the First National Conference on Artificial Intelligence, Stanford, Cal.

[11] Haberlandt, K. and Bingham, G. (1978). Verbs contribute to the coherence of brief narratives: Reading related and unrelated sentence triples. Journal of Verbal Learning and Behavior, 17, P-419-425.

[12] Rothkopf, E.Z. and Bisbicos, E. (1967). Selective facilitative effects of interspersed questions on learning from written material. Journal of Educational Psychology. 58. p. 56-61.

[13] Schank, R. C. (1975). Conceptual Information Processing. North-Holland, Amsterdam.

[14] Schank, R. C. and Abelson, R. P. (1977). Scripts, Plans, Goals and Understanding Erlbaum Press, Hillsdale, N.J.

[15] Sridharan,N.S. and Schmidt,C.F. (1978). Knowledge-directed inference in BELIEVER. In Waterman & Hayes-Roth (Eds.), Pattern-Directed Inference Systems. Academic Press, N.Y.

[16] Wilensky, R. (1978). Understanding Goal-based Stories. Ph.D. Thesis, Yale University, New Haven, Conn.