

SHAPING EXPLANATIONS:

Effects of questioning on text interpretation

Richard H. Granger, Jr.
Artificial Intelligence Project
Computer Science Department
University of California
Irvine, California 92717

ABSTRACT

Results in cognitive psychology have shown that readers can be steered away from an otherwise plausible interpretation of a story by extra-textual factors such as the source of the text, the stated reading purpose, interruptions and repetition of questions about the text. For instance, successive repetitions of the same question about a given text will often elicit a series of alternative interpretations of the text. This effect cannot be accounted for by established principles of text processing behavior, such as people's preference for cohesive and parsimonious representations of text. This paper presents a computer program called MACARTHUR, which models this behavior by varying 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. In light of the results, some new experiments are suggested in hopes of arriving at a new principle, beyond cohesion and parsimony, to account for the observed text processing behavior.

1.0 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 just wanted to avoid reporters on his vacation.

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. However, readers can be steered away from this explanation by external factors such as repetition of the same question. 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 inferences 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 text understanding systems (e.g., Cullingford [1978], Wilensky [1978], DeJong [1979], Charniak [1978]) do not account for people's ability to make different inferences depending on external factors such as re-probing. MACARTHUR's ability to re-direct its own inferences arises from a new classification scheme for explanations based on an attribute termed the "shape" of an explanation. The program 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, the concluding section of this paper proposes some possible new experiments, and some possible extensions to MACARTHUR.

2.0 BACKGROUND: COHESION AND PARSIMONY

2.1 The cohesion principle

Results in cognitive psychology have shown that people almost universally construct interpretations of text which serve to coherently connect the separate statements in a text, even when such connections are not at all obvious. For instance, Haberlandt and Bingham [1978] have found evidence for causal connective inferences being made among the sentences in examples like the following:

- [2] Brian punched George. George called the doctor. The doctor arrived.

[3] Brian punched George. George liked the doctor. The doctor arrived.

Subjects took longer to read [3] than [2], presumably spending the extra time trying to infer causal or intentional connective inferences among the statements in the text.

Similarly, Bower, Black and Turner [1979] found that reading times were longer when readers had to perform more than one inferential "step" to establish a causal connection between two statements in a narrative. These results and others providing evidence for spatial, instrumental, referential, causal and intentional connective inference have demonstrated that a crucial feature of human text understanding is the ability to construct a connected and coherent representation of a text. Taken together, these results form what we may term the "cohesion principle" of text processing behavior.

Researchers in AI have constructed a number of process models of text understanding which are consistent with the cohesion principle. AI programs that have addressed the problem of connectedness in texts include the MARGIE program (Schank [1975]) in terms of causal connections, the SAM and Ms.Malaprop programs in terms of script- and frame-based connections (Schank et al [1975], Cullingford [1978], Charniak [1979]), and the PAM and BELIEVER programs in terms of intentional connections (Schank and Abelson [1977], Wilensky [1978], Sridharan and Schmidt [1978]).

2.2 The parsimony principle

The cohesion principle alone is not sufficient to account for people's interpretations of text. For instance, consider the following deceptively simple example (from Granger [1980]):

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

When asked why Mary picked up the magazine, people in informal experiments overwhelmingly answer that she picked it up with the intention of swatting the fly. However, this answer corresponds to only one of (at least) three possible interpretations of the text, none of which can be ruled out on grounds of logic or the cohesion principle:

(4a) 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.

(4b) 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.

(4c) Mary picked up a magazine to swat a fly with it.

This same phenomenon occurs in any "garden path" text; i.e., a text that suggests an initially plausible inference which turns out to be "supplanted" (Granger [1980]) in the final rep-

resentation. To account for these observations, Granger proposed the Parsimony Principle, which states that the preferred interpretation of a text is the one in which the fewest number of inferred intentions of a story character account for the maximum number of his actions. This principle has been incorporated into a computer program called ARTHUR (A Reader THAT Understands Reflectively), which can supplant its own initial inferences in light of subsequent information in a text, thereby enabling it to read garden path stories.

3.0 THE SHAPE OF EXPLANATIONS

The cohesion and parsimony principles together still fail to account deterministically for certain text understanding behavior. In particular, people's ability to generate alternative interpretations of a text in response to re-questioning cannot be explained by these principles, since, for example, all four of the interpretations given earlier in this paper for story [1] are coherent and parsimonious.

In order to account for this behavior, we have developed a classification scheme for alternative explanations based on an attribute of explanations we term their "shape". This scheme has proven useful in the explanation-selection algorithm used by MACARTHUR in generating alternative interpretations of a text. Following is a list of the four 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 reacting; 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 which 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 recognize deception via conflicting beliefs held by different characters.

Following is an illustration of 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 previously given for this story can now be categorized 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)

Consider story [5], another story that MACARTHUR can process (see Granger [1981] for examples of detailed output from MACARTHUR):

- [5] Dr. Fitzsimmons 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.

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 reason for going in there, so he yawned to pretend he was tired. (COVER-STORY)

4.0 CONCLUSIONS: PROPOSED EXPERIMENTS

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 story, 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 monotonically 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 based on 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, non-straightforward texts like [1] and [5] in this paper could be used. According to the cohesion principle, 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 these non-straightforward texts. However, since there are a number of different alternative interpretations for these texts, 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 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.

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, 11, 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). Script Application: Computer Understanding of Newspaper Stories. Ph.D. Thesis, Yale University, New Haven, Conn.
- [7] Frederiksen, C.H. (1975). Effects of context-induced processing operations on semantic information acquired from discourse. Cognitive Psychology, 7, p. 371-458.
- [8] Granger, R. (1980). When Expectation Fails: Towards a self-correcting inference system. Proceedings of the First National Conference on Artificial Intelligence, Stanford, Cal.
- [9] Granger, R. (1981). Directing and Re-directing Inference Pursuit: Extra-Textual Effects on Text Interpretation. Technical Report #171, Computer Science Dept., University of California, Irvine, Cal. Submitted to the 7th IJCAI, Vancouver.
- [10] 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.
- [11] 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.
- [12] Schank, R. C. (1975). Conceptual Information Processing, North-Holland, Amsterdam.
- [13] Schank, R. C. and Abelson, R. P. (1977). Scripts, Plans, Goals and Understanding. Erlbaum Press, Hillsdale, N.J.
- [14] 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.
- [15] Wilensky, R. (1978). Understanding Goal-based Stories. Ph.D. Thesis, Yale University, New Haven, Conn.