

WHEN EXPECTATION FAILS:

Towards a self-correcting inference system

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

ABSTRACT

Contextual understanding depends on a reader's ability to correctly infer a context within which to interpret the events in a story. This "context-selection problem" has traditionally been expressed in terms of heuristics for making the correct initial selection of a story context. This paper presents a view of context selection as an ongoing process spread throughout the understanding process. This view requires that the understander be capable of recognizing and correcting erroneous initial context inferences. A computer program called ARTHUR is described, which selects the correct context for a story by dynamically re-evaluating its own initial inferences in light of subsequent information in a story.

INTRODUCTION

Consider the following simple story:

- [1] Geoffrey Huggins walked into the Roger Sherman movie theater. He 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 by most readers in informal experiments. On the basis of pure logic, either of these inferences is equally plausible. However, people overwhelmingly choose the first inference to explain this story, maintaining that Geoff did not go into the theater to see a movie.

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 about Geoff's goal in this story requires rejection of this initial inference. This paper describes a program called ARTHUR (A Reader That Understands Reflectively) which understands stories like [1] by generating tentative initial context inferences and then re-evaluating its own inferences in light of subsequent information in the story. By this process ARTHUR understands misleading and surprising stories, and expresses its surprise in English. For example, from the above story,

ARTHUR answers the following question about Geoff's intentions:

Q) Why did Geoff go into the movie theater?

A) AT FIRST I THOUGHT IT WAS BECAUSE HE WANTED TO WATCH A MOVIE, BUT ACTUALLY IT'S BECAUSE HE WANTED TO BUY COCAINE.

(For a much more complete description of ARTHUR, see Granger [1980]).

We call the problem of finding the correct inference in a story the "context-selection problem" (after the "script-selection problem" in Cullingford [1978] and Dejong [1979], which is a special case (see Section 4.2)). All the "contexts" (or "context inferences") referred to in this paper are goals, plans or scripts, as presented by Schank and Abelson [1977]. Other theories of contextual understanding (Charniak [1972], Schank [1973], Wilks [1975], Schank and Abelson [1977], Cullingford [1978], Wilensky [1978]) involve the selection of a context which is then used to interpret subsequent events in the story, but these theories fail to understand stories such as [1], in which the initially selected context turns out to be incorrect. ARTHUR operates by maintaining an "inference-fate graph", containing the tentative inferences generated during story processing, along with information about the current status of each inference.

BACKGROUND: SCRIPTS, PLANS, GOALS & UNDERSTANDING

2.1 Contextual understanding

ARTHUR's representational scheme is adopted from Schank and Abelson's [1977] framework for representing human intentions (goals) and methods of achieving those goals (plans and scripts). The problem ARTHUR addresses is the process by which a given story representation is generated from a story. It will be seen that this process of mapping a story onto a representation is not straightforward, and may involve the generation of a number of intermediate representations which are discarded by the time the final story representation is complete.

Recall the first sentence of story [1]: "Geoffrey Huggins walked into the Roger Sherman movie theater." ARTHUR's attempt to infer a context for this event is based on knowledge of typical functions associated with objects and locations. In this instance, a movie theater is a known location with an associated "scripty" activity: viewing a movie. Hence, whenever a story character goes to such a location, one of the plausible inferences from this action is that the character may intend to perform this activity. Seeing a movie also has a default goal associated with it: being entertained. Thus, ARTHUR infers that Geoff plans to see a movie to entertain himself.

When the next sentence is read, "He went up to the balcony, where Willy North was waiting with a gram of cocaine", ARTHUR again performs this bottom-up inference process, resulting in the inference that Geoff may have been planning to take part in a drug sale. Now ARTHUR attempts to connect this inference with the previously inferred goal of watching a movie for entertainment. Now, however, ARTHUR fails to find a connection between the goal of wanting to see a movie and the action of meeting a cocaine dealer. Understanding the story requires ARTHUR to resolve this connection failure.

2.2 Correcting an erroneous inference

Having failed to specify a connecting inferential path between the initial goal and the new action, ARTHUR generates an alternative goal inference from the action. In this case, the new inference is that Geoff wanted to entertain himself by intoxicating himself with cocaine. (Note that this inference too is only a tentative inference, and could be supplanted if it failed to account for the other events in the story.) ARTHUR now has a disconnected representation for the story so far: it has generated two separate goal inferences to explain Geoff's two actions. ARTHUR thinks that Geoff went to the theater in order to see a movie, but that he then met up with Willy in order to buy some coke. This is not an adequate representation for the story at this point. The correct representation would indicate that Geoff performed both of his actions in service of a single goal of getting coke, and that he never intended to see a movie there at all; the theater was just a meeting place.

Hence, ARTHUR instead infers that Geoff's action of going into the theater was in service of the newly inferred goal, and discards the initial inference (wanting to see a movie) which previously explained this action. We call this process supplanting an inference: ARTHUR supplants its initial "see-movie" inference by the new "yet-coke" inference, as the explanation for Geoff's two actions.

ARTHUR's representation of the story now consists of a single inference about Geoff's intentions (he wanted to acquire some coke) and two plans performed in service of that goal (getting to the movie theater and getting to Willy), each of which was carried out by a physical action (PTRANSing to the theater and PTRANSing to Willy). At this point, the initial goal inference (that Geoff wanted to see a movie) has been supplanted: it is no longer considered to be a valid inference about Geoff's intentions in light of the events in the story.

OPERATION OF THE ARTHUR PROGRAM

3.1 Annotated run-time output

The following represents actual annotated run-time output of the ARTHUR program. The input to the program is the following deceptively simple story:

[2] 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. (ARTHUR's output has been shortened and simplified here for pedagogical and financial reasons.)

: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 underlying 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].

3.2 The parsimony principle

ARTHUR's answer as given above is not the only possible interpretation of the story; it is only one of the following three alternatives, all of which are valid on the basis of what the story says:

- (2a) 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.
- (2b) 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.
- (2c) Mary picked up a magazine to swat a fly with it.

The last interpretation (2c) 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.

In [2], as in [1], the interpretation generated by the reader is the most parsimonious of the possible interpretations. That is, the preferred interpretation is the one in which the fewest number of inferred goals of a story character account for the maximum number of his actions. We summarize this observation in the following principle:

The Parsimony Principle

The best context inference is the one which accounts for the most actions of a story character.

In other words, the decision to replace a previous inference by a new one is not based on the explicit contradiction of that inference by subsequent information in the story. Example [2], for instance, has three possible interpretations, none of which can be ruled out on strictly logical grounds. Rather, the reader prefers the most parsimonious story representation over less parsimonious ones, that is, the representation which includes the fewest goal inferences to account for the actions in the story. This is true even when it requires the reader to do the extra work of replacing one of its own previous inferences, as in example [2].

CATEGORIES OF ERRONEOUS INFERENCES

4.1 Goals

ARTHUR is capable of recognizing and correcting erroneous context inferences in order to maintain a parsimonious explanation of a story. The examples given so far have dealt only with erroneous goal inferences, but other conceptual categories of inferences can be generated erroneously as well. In this section, examples of other classes of erroneous inferences will be given, and it will be shown why each different class presents its own unique difficulties to ARTHUR's correction processes.

4.2 Plans and scripts

Consider the following simple story:

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

This is an example in which ARTHUR correctly infers the goal of the story character, but erroneously infers the plan that he is going to perform in service of his goal. ARTHUR first infers that Carl will read the newspaper to alleviate his boredom, but this inference fails to explain why Carl then gets his tennis racket. ARTHUR at this point attempts to supplant the initial goal inference, but in this case ARTHUR knows that that goal was correctly inferred, because it was implicitly stated in the first sentence of the story (that Carl was bored). Hence ARTHUR infers instead that it erroneously inferred the plan by which Carl intended to satisfy his goal (reading the newspaper). Rather, Carl planned to alleviate his boredom by playing tennis.

The problem now is to connect Carl's action of picking up the newspaper with his plan of playing tennis. Instead of using the newspaper as a functional object (in this case, reading material), Carl has treated it as an instrumental object that must be moved as a precondition to the implementation of his intended plan. (Preconditions are discussed in Schank and Abelson [1977]). ARTHUR recognizes that an object can be

used either functionally or instrumentally. Furthermore, when an action is performed as a precondition to a plan, typically the objects used in the action are used instrumentally, as in [3]. ARTHUR's initial inference about Carl's plan was based on the functionality of a newspaper. It is able to supplant this inference by an inference that Carl instead used the newspaper instrumentally, as a precondition to getting to his tennis racket, which in turn was a presumed precondition to using the racket to play tennis with. Hence, correcting this erroneous plan inference required ARTHUR to re-evaluate its inference about the intended use of a functional object.

4.3 Causal state changes

Consider the following example:

[4] Kathy and Chris were playing golf. Kathy hit a shot deep into the rough.

We assume that Kathy did not intend to hit her ball into the rough, since she's playing golf, which implies that she probably has a goal of winning or at least playing well. Her action, therefore, is probably not goal-oriented behavior, but is accidental: that is, it is an action which causally results in some state which may have an effect on her goal.

This situation differs from stories like [1], [2] and [3], in that ARTHUR does not change its mind about its inference of Kathy's goal. Rather than assuming that the goal inference was erroneous, ARTHUR infers that the causal result state hinders the achievement of Kathy's goal. Any causal state which affects a character's goal, either positively or negatively, appears in ARTHUR's story representation in one of the following four relationships to an existing goal:

- 1 - the state helps the achievement of the goal;
- 2 - the state hinders achievement of the goal;
- 3 - the state achieves the goal entirely; or
- 4 - the state thwarts the goal entirely.

If ARTHUR did assume that Kathy's shot was intentional, then the concomitant inference is that she didn't really want to win the game at all; or, in other words, that the initial inference was erroneous. This is the case in the following example:

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

Understanding this story requires ARTHUR first to infer that Kathy intends to win the game; then to notice that her action has hindered her goal, and finally to recognize that the initial goal inference was erroneous, and to supplant it by the inference that Kathy actually intended to lose the game, not win it.

4.4 Travelling down the garden path ...

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

Since ARTHUR operates by generating tentative initial inferences and then re-evaluating those inferences in light of subsequent information, ARTHUR understands simple garden path stories. Not all garden path stories cause us to experience surprise. For example, many readers of story [2] do not notice that Mary might have been planning to read the magazine, unless that intermediate inference is pointed out to them. Hence we hypothesize that the processes involved in understanding stories with surprise endings must differ from the processes of understanding other garden path stories. Hence, ARTHUR's understanding mechanism is not entirely psychologically plausible in that it does not differentiate between stories with surprise endings and other garden path stories.

A more sophisticated version of ARTHUR (call it "Macro-ARTHUR") might differentiate between "strong" default inferences and "weak" tentative inferences when generating an initial context inference. If a strong initial inference is generated, then MacARTHUR would consciously "notice" this inference being supplanted, thereby experiencing surprise that the inference was incorrect. Conversely, if the initial inference is weak, MacARTHUR may not commit itself to that inference, but rather may choose to keep around other possible alternatives. In this case MacARTHUR would only experience further specification of the initial tentative set of inferences, rather than supplanting a single strong inference. The question of when readers processes consciously versus unconsciously is still an open question in psychology. Future psychological studies of the cognitive phenomena underlying human story understanding (such as in Thorndyke [1976], [1977], Norman and Bobrow [1975], and Hayes-Roth [1977], to name a few) may be able to provide data which will shed further light on this issue.

5.1 Process and representation in understanding

This paper has presented a process for building story representations which contain inferences not explicitly stated in the story. The representations themselves are not new; they are based on those presented by Schank and Abelson [1977]. What is new here is the process of arriving at a story representation. Most contextual understanders (e.g. Charniak [1975], Cullingford [1978], Wilensky [1978]) would fail to arrive at the correct story representations for any of the examples in this paper, because initial statements in the examples trigger inferences which prove to be erroneous in light of subsequent story statements. ARTHUR's processing of these examples shows that arriving at a given story representation may require the reader to generate a number of intermediate inferences which get discarded along the way, and which therefore play no role in the final representation of the story.

Thus a final story representation may not completely specify the process by which it was generated, since there may have been intermediate inferences which are not contained in the final representation. Yet we know that when people have understood one of these examples, they can express these intermediate inferences with phrases like "At first I thought X, but actually it's Y." ARTHUR keeps track of its intermediate inferences while understanding a story, and maintains an "inference-fate graph" containing all inferences generated during story processing, whether they end up in the final story representation or not.

The point here is that the relationship between a given story representation on the one hand, and the process of arriving at that representation on the other, may be far from straightforward. The path to a final story representation may involve sidetracks and spurious inferences which must be recognized and corrected. Therefore, specifying the representations corresponding to natural language inputs is not enough for a theory of natural language processing; such a theory must also include descriptions of the processes by which a final representation is constructed. ARTHUR has demonstrated one area in which specification of process and representation diverge: the area of correcting erroneous inferences during understanding. Further work will be directed towards specifying other conditions under which process and representation are not straightforwardly related in natural language tasks.

- [1] Cullingford, R. (1978). Script Application: Computer Understanding of Newspaper Stories. Ph.D. Thesis, Yale University, New Haven, Conn.
- [2] Granger, R. (1980). Adaptive Understanding: Correcting Erroneous Inferences. Ph.D. Thesis, Yale University, New Haven, Conn.
- [3] Hayes-Roth, B. (1977). Implications of human pattern processing for the design of artificial knowledge systems. In Pattern-directed inference systems (Waterman and Hayes-Roth, eds.). Academic Press, N.Y.
- [4] Kintsch, W. (1977). Memory and Cognition. Wiley, New York.
- [5] Norman, D. and Bobrow, D (1975). On the role of active memory processes in perception and cognition. In The structure of human memory (Coffey, C., ed.). Freeman, San Francisco.
- [6] Schank, R. C. and Abelson, R. P. (1977). Scripts, Plans, Goals and Understanding. Erlbaum Press, Hillsdale, N.J.
- [7] Thorndyke, P. (1977). Pattern-directed processing of knowledge from texts. In Pattern-directed inference systems (Waterman and Hayes-Roth, eds.). Academic Press, N.Y.
- [8] Wilensky, R. (1978). Understanding Goal-based Stories. Ph.D. Thesis, Yale University, New Haven, Conn.
- [9] Wilks, Y. (1975). Seven Theses on Artificial Intelligence and Natural Language, Research Report No. 17, Istituto per gli Studi Semantici e Cognitivi, Castagnola, Switzerland.