REASONS FOR BELIEFS IN UNDERSTANDING: APPLICATIONS OF NON-MONOTONIC DEPENDENCIES TO STORY PROCESSING

Paul O'Rorke

Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL 61801

ABSTRACT

Many of the inferences and decisions which contribute to understanding involve fallible assumptions. When these assumptions are undermined, computational models of comprehension should respond rationally. This paper crossbreeds AI research on problem solving and understanding to produce a hybrid model ("reasoned understanding"). In particular, the paper shows how non-monotonic dependencies [Doyle79] enable a schema-based story processor to adjust to new information requiring the retraction of assumptions.

I INTRODUCTION

Many of the inferences and decisions involved in understanding "jump to conclusions" which might later turn out to be erroneous. For example, upon reading "John put a quarter in the slot", a video game addict may jump to the conclusion that John is playing a video game. If the addict is then told "John put another quarter in and pushed a button for a cola", he should revise his beliefs.

How can a computational model of understanding adjust efficiently to new information which invalidates previous assumptions?

The solution proposed in this paper applies a view of rational thought developed in AI research on problem solving [Doyle79] [de Kleer79] [Charniak80] to current schema-based models of understanding. The resulting hybrid view of comprehension will be referred to as "reasoned understanding". In this view, comprehension is the work of rules which infer and justify new beliefs on a basis of old beliefs (many of which are compactly specified with the aid of schemata). Justifications of beliefs record their dependence on inference rules and other beliefs. When fallible assumptions lead to conflicting beliefs these dependencies may be used to determine and revise the incompatible assumptions underlying the conflict.

This report describes work done in the AI group of the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign. It was supported in part by National Science Foundation Grant IST 81-20254 and by the Air Force under grant F49620-82-K-0009.

This paper describes the design and operation of a schema-based story processor called RESUND. The design applies non-monotonic dependencies and associated processes (like dependency-directed backtracking [Doyle79]) to solve some basic belief revision problems which arise in natural language processing.

Consider the following example from [Collins80].

1. "He plunked down $5 at the window."
2. "She tried to give him $2.50, but he refused to take it."
3. "So when they got inside, she bought him a large bag of popcorn."

People commonly make two interesting mistakes on reading this window text:

The first mistake is made by people who assume after the first sentence that "he" is making a 5$ bet at a horse race. In spite of this, one usually concludes with the assumption "they" are going to see a movie. How might a story processor retract its presumption of one scenario (eg. "going to a horse race") in favor of another ("going to the movies")?

The second common mistake is the initial identification of the second sentence as an attempt to return change, in spite of "his" disturbing refusal to accept it. Later, we see that "she" was not the attendant at the window, but that "she" is accompanying "him". Her attempt to give him 2.50$ is re-explained as an attempt to pay her own way. This new interpretation of her action makes his refusal comprehensible. How might a story processor recover from misidentification of events and objects?

Section three shows how a schema-based story processor can recover from these mistakes using the techniques described in section two.

II DEPENDENCIES IN UNDERSTANDING

Reasoned understanding is a view of comprehension inspired by AI research on problem solving and decision making [Doyle80]. It inherits the notion that attitudes (eg. beliefs) are the important indicators of the state of comprehension. In fact, understanding is viewed as a process of transition from one set of "current" attitudes to another. Furthermore, justifications for attitudes play a key role in determining which ones are currently held. For example, reasons supporting possible beliefs not only determine which ones are currently believed, but also provide the basis for inferring new beliefs and retracting old ones. This section sketches the design and operation of a schema-based story processor (called RESUND) compatible with reasoned understanding. The story processor is still under construction at the time of this writing.

The "knowledge base" of the story processor is organized into bundles of assertions called schemata. A schema intended to capture knowledge about an event includes variables for objects which play roles in the event and a list of associated schemata. Relationships between the associated schemata include primitive temporal and causal links. Top down script elaboration inferences (a la SAM) [Cullingford81] are supported by a special elaboration relation which specifies the consequences of belief in an assertion that an event has occurred. Construction of intentional explanations (a la PAM) [Wilensky81] is supported by associations such as X is an aim (goal) of Y and X is a method for achieving Y. In addition, schemata specify constraints and default values for their role variables.

A crucial point about inferences in understanding is that they are often presumptive. They generate beliefs not just by logical deduction, but also by making assumptions of various kinds. The main difference between RESUND and previous story processors is that RESUND uses non-monotonic dependencies to compensate for the fact that its inference processes are fallible. A collection of inference rules (a la AMORD) [de Kleer79] generate and justify new assertions representing beliefs. Some inference mechanisms which contribute to the construction of explanations of sequences of input events are elaboration, intentionality, identification and criteriality.

Elaboration rules expand the definitions of complex concepts in schemata to capture inferences similar to SAM's [Cullingford81] top down script applications. For example, when RESUND "comes to believe" that a complex event has taken place, (eg. John purchased some tickets from someone), the definitional consequences (John paid someone, John recieved some tickets from someone) are inferred by elaboration. Intentionality rules generate intentional explanations similar to those constructed by PAM [Wilensky81] using the information about goals and methods supplied by schemata. When two descriptions appear to refer to the same event or object, RESUND's identification inference rules make the assumption that the descriptions are co-referential. Bottom up schema invocation is also treated as a presumptive inference in RESUND (called criteriality). Event induced schema activation is the simplest kind of criteriality assumption. This happens when a sub-event of a complex event causes RESUND to assume the complex event occured (as when "She bought popcorn" suggests "she went to the movies") [Schank77] [DeJong79].

AI natural language processing systems are bound to make mistakes like those made by people reading the window text. RESUND recognizes mistakes which reveal themselves in the confusion of conflicting beliefs. The most common conflicts in the examples studied to date have been identity conflicts and other violations of schematic constraints or defaults. Identity conflicts arise when two objects or events are assumed to be identical and not identical simultaneously. Schematic constraint violations occur when a restriction on variables or other schemata associated with a given schema is broken. When such conflicts arise, a process like dependency-directed backtracking determines the underlying incompatible assumptions by looking back along dependencies. Unfortunately, in natural language processing, it wouldn't do to just arbitrarily rule out one of these assumptions in order to resolve the conflict. Thus, RESUND requires a method to decide which assumption should be revoked (which assumption is weakest).

The current design calls for a collection of preference policies which represent different criteria for gauging the relative strength of incompatible assumptions. Some of these policies implement text comprehension principles and problem solving strategies reported in [Collins80] and [Wilensky83]. Another class of policies is based on RESUND's taxonomy of inferences and the notion that some assumptions made in natural language processing are inherently weaker (more fallible) than others. These policies prefer constraints and elaboration inferences over defaults and see identification, intentionality, and criteriality assumptions as the most likely losers in a conflict. Unfortunately, several preference policies may be applicable to a given conflict. The current design calls for the simplest possible conflict resolution: a total order on the preference policies. When several are applicable, the strongest one is allowed to choose the weakest assumption.

III EXAMPLES

To see how RESUND will handle mis-identification, consider the following simplified variant of the window text.

Mis-identification Example.

1. He put down $5 at the Thunderbird theatre ticket window.
2. She gave him $2.50.
3. When they got inside, she bought herself a bag of popcorn.

The first sentence invokes a schema about "going to the movies." Elaboration of this action includes "purchasing tickets," which includes "paying for the tickets," and the "return of change" (if any). The placement of $5 at the ticket window is identified as part of "paying" for some tickets. This means "he" is the BUYER of tickets, a member of the PARTY "going to the movies". By convention, the roles in this and other schemata will be in upper case.

The action in the second sentence is identified as the "return of change". This isn't the only possible identification of the action, because there are several other transfers of money associated with "going to the movies". None of the other identifications is compatible with this one, and at this point "return of change" is preferred because it is contained in the schema describing "buying tickets", as was the initial transfer of payment. This action identification implies two new role identifications: she is the TICKET-BOOTH-ATTENDANT and the $2.50 is his CHANGE.

Both actions in the third sentence are identified as actions of members of the PARTY going to the movie, (namely "entering the theater" and "purchasing refreshments"). Thus, "she" is seen to be a moviegoer. This violates a schematic constraint on TICKET-BOOTH-ATTENDANT which reflects the fact that normally, the attendant is not a member of one's party when one goes to the movies.

Whether this constraint is learned by experience or derived from more basic constraints, recognition of the constraint violation triggers dependency-directed backtracking. Data-dependencies associated with the inconsistent beliefs (TICKET-BOOTH-ATTENDANT in PARTY, not TICKET-BOOTH-ATTENDANT in PARTY) and their ancestors are examined to determine the assumptions underlying the conflict. The belief that the ATTENDANT should not be in PARTY is strong because it is based on constraints in the "going to the movies" schema. The belief that ATTENDANT is in PARTY depends on she = ATTENDANT and she in PARTY. Ultimately, the conflict depends on the following identification assumptions.

A1. The purchase of refreshments by a member of the PARTY attending the movie = "she bought popcorn."
A2. The return of change by the TICKET-BOOTH-ATTENDANT = "she gave $2.50."

There is no alternative to identification A1 in the "going to the movies" schema, but A2 does have an alternative: members of PARTY sometimes repay (or prepay) the BUYER who purchases the group's tickets. For this reason, A2 is considered weaker than A1, and the constraint violation conflict is resolved by ruling A2 out.

Next, consider an example of the problem of mis-activation of schemata.

Mis-activation Example.

1. John put two quarters in the slot.
2. Then he started his first game.

If you originally thought John was going to get a cola from a vending machine, as he did earlier in the paper, then you had to retract this assumption with any conclusions founded on it. This example, like the horse-races to movies switch, is just about the simplest kind of mis-activation.

We have worked out dependencies for event induced schema activation which enable RESUND to recover from this sort of mistake:

The insertion of quarters in a slot is an action which invokes the "coke machine" and "video game" schemata by event induced activation [DeJong79] Elaborations of these schemata include corresponding insertions of change, as well as inferences about what will happen next, etc. An identity conflict arises because the insertion of change in the coke machine schema is incompatible with the insertion in the video game schema, (if they are not compatible, they cannot both be identical to the input event). Dependency-directed backtracking determines that one of the schema invocations must be retracted. A preference policy decides to retract video-game on very weak grounds, perhaps because the last time the system saw this sentence, "coke machine" turned out to be the right schema.

The second sentence contains an event which is identified as part of the dead "video-game" schema. The fact that there is no alternative identification is seen as an argument against the weak decision to rule out the video game scenario. The original contradiction comes back in, but now there is a strong reason to prefer "video-game" over "coke-machine": namely that it explains more input events.

IV CONCLUSION

This paper argues that better models of understanding can be constructed by applying views and techniques originally developed in AI research on problem solving. In particular, nearly all current story processors have no reasons for their beliefs; so when they make inferences and decisions which jump to false conclusions, they have no recourse to reasoned retraction. ARTHUR, MCARTHUR, and JUDGE are exceptional, in that they attempt to recover from inference errors [Granger80] [Granger82]. However, they appear to concentrate exclusively on revision of intentional explanations, and do not appear to use dependency-network-maintenance techniques to supplant incorrect explanations.

The usefulness of reasoned understanding as part of a model of human comprehension is limited by the fact that it ignores affect and emotion. In addition, we have only described a handful of belief revision methods and they are so simple that people usually do them unconsciously.

Nevertheless, we expect non-monotonic dependency networks and associated processes (like dependency-directed backtracking) to become integral parts of future applied AI natural language processing systems. In an effort to help make this happen, we have begun implementing RESUND. The design incorporates schemata and an inference taxonomy derived from a schema-based story processor constructed in the summer of 1982. We have worked out detailed dependencies and preference policies which seem necessary for several examples like those presented in section III. When the implementation is complete we will run experiments to verify that the design works as planned. We expect future experiments to lead to the discovery of new types of inference and new preference policies, if not to radical changes in the design of the story processor.

ACKNOWLEDGEMENTS

My sincere thanks to G. DeJong, M. Dorfman, C. Debrunner and the other reviewers of this paper.

REFERENCES

[Charniak80] E. Charniak, C. Riesbeck and D. McDermott, "Data Dependencies," in Artificial Intelligence Programming, Lawrence Erlbaum Associates, Hillsdale, N.J., 1980, 193-226.

[Collins80] A. Collins, J. S. Brown and K. M. Larkin, "Inference in Text Understanding," in Theoretical Issues in Reading Comprehension, R.J. Spiro, B.C. Bruce, and W.F. Brewer (ed.), Lawrence Erlbaum Associates, Hillsdale, N.J., 1980, 385-407.

[Cullingford81] R. Cullingford, "SAM," in Inside Computer Understanding: Five Programs Plus Miniatures, R.C. Schank and C.K. Riesbeck (ed.), Lawrence Erlbaum Associates, Hillsdale, N.J., 1981, 75-119.

[de Kleer79] J. de Kleer, J. Doyle, G. L. Steele and G. J. Sussman, "Explicit Control of Reasoning," in Artificial Intelligence: An MIT Perspective, vol. 1, P.H. Winston and R.H. Brown (ed.), MIT Press, Cambridge, Massachussetts, 1979, 33-92.

[DeJong79] G. F. DeJong, "Skimming Stories in Real Time: An Experiment in Integrated Understanding," 158, Yale University Dept. of Comp. Sci., New Haven, Conn., May 1979.

[Doyle79] J. Doyle, "A Truth Maintenance System," Artificial Intelligence 12, (1979), 231-272.

[Doyle80] J. Doyle, "A Model for Deliberation, Action, and Introspection," Artificial Intelligence Technical Report-581, MIT Artificial Intelligence Lab., Cambridge, MA, May, 1980.

[Granger80] R. H. Granger, "Adaptive Understanding: Correcting Erroneous Inferences," 171, Yale Univ. Dept. of Comp. Sci., New Haven, Conn., January 1980.

[Granger82] R. H. Granger, "Judgemental Inference: A Theory of Inferential Decision-Making During Understanding," Proc. of the Fourth Annual Conf. of the Cognitive Science Society, Ann Arbor, MI, 1982, 177-180.

[Schank77] R. Schank and R. Abelson, Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures, Lawrence Erlbaum, Hillside, NJ, 1977.

[Wilensky81] R. Wilensky, "PAM," in Inside Computer Understanding: Five Programs Plus Miniatures, R.C. Schank and C.K. Riesbeck (ed.), Lawrence Erlbaum Associates, Hillsdale, N.J., 1981, 136-179.

[Wilensky83] R. Wilensky, Planning and Understanding: A Computational Approach to Human Reasoning, Addison-Wesley, Reading, Mass., 1983.