CONVEYING IMPLICIT CONTENT IN NARRATIVE SUMMARIES

Malcolm E. Cook, Wendy G. Lehnert, David D. McDonald

Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts 01003

ABSTRACT

One of the key characteristics of any summary is that it must be concise. To achieve this the content of the summary (1) must be focused on the key events, and (2) should leave out any information that the audience can infer on their own. We have recently begun a project on summarizing simple narrative stories. In our approach, we assume that the focus of the story has already been determined and is explicitly given in the story's long-term representation; we concentrate instead on how one can plan what inferences an audience will be able to make when they read a summary. Our conclusion is that one should think about inferences as following from the audience's recognition of the central concepts in the story's plot, and then plan the textual structure of the summary so as to reinforce that recognition.

BACKGROUND

This research builds on our previous work on narrative structure and generation. We are using Plot Units [Lehnert 1981] to represent the structure of the original narrative, and use Mumble [McDonald 1983] to do the linguistic realization. To connect these two facilities we have a new interface and a new text planning component named "Precis".

Plot units are a technique for organizing the conceptual representation of a narrative in such a way that the topological structure of the representation directly indicates which events are central to the story and which are peripheral. A graph of connected plot units is constructed for a story as it is understood, based on the recognition of goal-oriented behavior by the characters and their affective reactions to events. Plot units summarize larger-scale relationships among explicit and implicit events in the story, and are oriented toward long term recall rather than appreciation of story style or specific wording.

Mumble is a "realization" module for language generation; it takes a stream of output from a text planner and incrementally produces fluent, cohesive English text in accordance with the planner's specifications. The planner decides what information should be imparted and most of its rhetorical features; Mumble filters those decisions in accordance with grammatical constraints, handles syntax and morphology, and performs the "smoothing" operations that are required by the discourse context in which the information appears.

1. This research was supported in part by the National Science Foundation under contracts IST-8217502 and IST-8104984, and in part by the Office of Naval Research under contract N00014-83-K-0580.

Precis stands between the plot unit graph and Mumble. It has been under development for a only short time and the ultimate form that its architecture will take is not yet fixed. We have so far been working bottom up, experimenting with different ways to combine the partial texts contributed by individual units and affect states, and trying to understand the consequences of the alternatives. We report here on one key "tactical" problem in narrative summarization which we refer to as *conceptual ellipsis*, omitting those events from a summary that we expect an audience to be able to infer on their own, and reinforcing that inference through a judicious choice of textual form.

THE NEED FOR CONCEPTUAL ELLIPSIS

Ever since the original work by Bartlett, researchers have appreciated that people who are remembering a story some time after they have heard it typically fail to distinguish between events that were explicitly stated in the story and those that they only inferred while reading it. Present day story understanding systems act in a similar way by maintaining only a single conceptual record of what they have understood regardless of its source [Joshi & Weischedel 1977, Graesser 1980, Dyer 1983]. Since our summarization process starts from the conceptual representation of the story rather than the text itself, it too will be unable to make this distinction.

This theory of memory has two consequences. One is that any decisions about what constituted the crux or point of the story must have been made at comprehension time rather than summarization time. This is one of the purposes of a plot unit representation. The other is that we now need to deliberately recalculate what information should be explicit in our summary and what should be left for the audience to infer; were this not done, the superfluous information in the summary would make it sound quite unnatural—as though it were being told by a person from a different society who did not have any commonsense understanding of the social context in which the story was set. How the explicit versus left-to-inference calculation turns out will vary with the summary: the same story can be summarized or retold in different ways depending on which character's point of view is taken or which events are emphasized. The plot unit graph is neutral on this question, and it will be an important part of what we do next in this research.

Decisions about conceptual ellipsis are made prior to any of the linguistic decisions about form; they are however linked to those later decisions since some linguistic forms will be more effective than others in indicating to the audience that an inference is intended. Certain marked choices of form will suggest to the reader that particular implications were "in the mind of the writer" at the time of generation. The conceptual decisions are thus the source

of dependencies that must be carried forward to the point where the text-form decisions will be made in order that the right realizations are chosen. By the same token there will also be dependencies percolating back to the conceptual ellipsis decisions indicating what alternative realizations are actually available in a given case and thus whether a particular implication can be adequately supported by the information that is included and the way it is phrased.

AN EXAMPLE

The following simple story will demonstrate the general phenomenon.

THE COMSYS STORY

John and Mike were competing for the same job at IBM. John got the job and Mike decided to start his own consulting firm, COMSYS. Within three years, COMSYS was flourishing. By that time, John had become dissatisfied with IBM so he asked Mike for a job. Mike spitefully turned him down.

A analysis of this text in terms of plot units has "Competition" as a central unit in the graph, which would make it a candidate basis for a summary of the story. All competition units have this pattern:

COMPETITION

Agent1	Agent2
M1	M2
+	-

Underlying this level of representation are the actual goals and events experienced by the two characters. In any competition unit, we have:

```
M1 : goal(agent1,goal1)
M2 : goal(agent2,goal2)
+ : success(goal1,event1)
- : failure(goal2,event2)
```

with the additional constraints:

```
C1 : event1 = event2
C2 : goal1 and goal2 cannot both be realized.
```

(Note that in C1 the positive and negative actualizations are actually the same event but from the point of view of two different characters.)

In the COMSYS story the competition is between John and Mike over who will get a particular job at IBM. The instantiation of the Competition unit in this story is:

```
M1 : A-goal1 (John has-role #employee in M-job1
              (where #employer = IBM))
M2 : A-goal2 (Mike has-role #employee in M-job1
              where #employer = IBM)
+ : success(A-goal1 , $hire(IBM,John))
- : failure(A-goal2 , not($hire(IBM,Mike)))
```

where

```
c1 : event1 = event2 = hire(IBM,John)
c2 : A-goal1 and A-goal2 cannot both be realized.
```

At the time of this writing, Precis can specify any of the following texts for this instantiation of the Competition unit, preferences dictated by conceptual ellipsis aside. (Discourse fluency effects such as verb phrase deletion or pronominalization are put in by Mumble as it is realizing Precis' specification.)

- (a) "John wanted to work for IBM and so did Mike. They hired John and did not hire Mike."
- (b) "Both John and Mike wanted to work for IBM, but they hired John."
- (c) "Mike wanted to work for IBM, but they hired John."

These three choices vary according to how much of the content of the Competition unit they explicitly express. Choice A includes each of the four affect states (M1, M2, +, -), smoothed somewhat by the recognition that M1 and M2 share the same predicate. The very simplest choice—one that did not express that commonality in its textual structure, e.g. "John wanted to work for IBM. Mike wanted to work for IBM. They hired John and did not hire Mike."—is completely unnatural; people wouldn't say it. This minimal level of implicit information that the textual structure must carry is accordingly not even made Precis' responsibility, but is instead carried out automatically within Mumble. The alternative realization of this commonality, using a conjoined subject rather than verb phrase deletion, is taken to be a stylistic decision and is not deliberated over by Precis.

If we begin to include the constraints that accompany the Competition unit (C1 and C2) explicitly in the summary then we can leave out more of the affect states as inferable. In choice B we make use of the first constraint, i.e. that the positive and the negative actualizations are consequences of the same event, to enable the omission of event2, not(hire(Mike,IBM)), from the text of the summary by dropping the phrase "they did not hire Mike".

In our present version of choice B there is no structural indicator of the constraint. It is probably no coincidence then that the text for B sounds a little odd—readers unfamiliar with the original story will not really understand what the *but* is supposed to be communicating until they go further and make the deduction that there must only have been one job available. A better version of B would probably be: "Both John and Mike wanted to work for IBM, but they *only* hired John", with the *only* acting as an explicit structural indicator of the information in the constraints. This addition can probably be licensed as a consequence of the second constraint that only one of the two goals can be realized. At the time of this writing we do not yet have an adequately general mechanism for making this observation and incorporating the *only*, so we have not included it among Precis' choices.

It is intriguing that choice c, "Mike wanted to work for IBM, but they hired John", is probably the best of the three choices even though it requires the audience to do the most inferencing. In c we have omitted state M1—that John wanted to work for IBM—yet the audience is able to recover this information quite easily given the presence of the *but*. Given the ease with which choice c is understood,

we are led to the suggestion that there may be a very general "template" being recognized here—that choice c is seen by an audience as an instance of the pattern:

<expression of agent A's goal>,
but <realization of agent B's goal>

and that this template always carries with it the inference that the two goals must be incompatible and therefore A's goal has not been satisfied.

Note that here again the choice would be improved by including an explicit lexical indication of the constraint: "*Mike wanted to work for IBM, but they hired John instead*". We expect that most instances of these "rhetorical markers" in texts will turn out to be indicators of constraint-level information akin to our present cases, which raises the intriguing possibility that a general theory of how they are used might arise out of this kind of work in generation.

SUMMARY

Currently, we are working with two programs. PUGG (Plot Unit Graph Generator) operates on an affect-state representation of a story, and produces a graph or network of plot units that act as pointers to the core of the conceptual representation of the input story and organizes how it will be "presented" to the program that plans the text of the summary, *Precis*. *Precis* is in the early stages of its development and so far can only use a single, core plot unit from the graph as the basis of the summary of the story.

Precis works at the interface between purely conceptual and purely linguistic considerations as it makes its planning decisions. It chooses from a set of alternative specifications for the summary that vary according to which of the elements of the plot unit are included and which left to be inferred by the audience once they recognize the story as a case of competition. *Precis* can state the three alternative choices described above (and a few other sets like them), and *Mumble* can take those specifications and produce the indicated texts. However we do not as yet have any general mechanism for deciding which choice to prefer over the others. Perhaps such a decision mechanism will become apparent once these single unit summaries are embedded in a larger context, or possibly there is no reasonable basis for decision without more knowledge of the purpose of the summary or the ability of a particular audience to make these kinds of inferences (one might have to talk quite differently to young children for example). In future work we also hope to be able to work out a specific basis for planning the use of inference-directing words like *only* or *instead*.

REFERENCES

- Dyer, M. (1983) *In-depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*, Cambridge, Mass.: MIT Press.
- Graesser, A. & C. (1981) *Prose Comprehension Beyond the Word* New York, N.Y.: Springer-Verlag.
- Joshi, A.K., and Weischedel (1977) Computation of a subclass of inferences: Presupposition and Entailment, in *Am J. of Comp. Linguistics*
- Lehnert, W. (1982) Plot Units: A Narrative Summarization Strategy, in Lehnert, W. and Ringle, M. (Eds.), *Strategies for Natural Language Processing*, Hillsdale, N.J.: Lawrence Erlbaum Associates.

Lehnert, W. (1983) "Narrative Complexity Based on Summarization Algorithms," *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Karlsruhe, Germany.

McDonald, D. (1983) "Natural Language Generation as a Computational Problem - an Introduction" in Brady, M. and Berwick, R. (Eds.) *Computational Models of Discourse*, Cambridge, Mass.: MIT Press.

McDonald, D. (1982) "Description Directed Control: its Implications for Natural Language Generation", in Cercone (ed.) *Computational Linguistics*, Dublin: Pergamon Press.