AN INTRODUCTION TO PLOT UNITS

Wendy G. Lehnert
Cynthia L. Loiselle

COINS Technical Report 84–13

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

July 1984

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

Abstract

Computers, like people, must use an internal memory representation in order to process narratives. By examining this representation we can learn more about how computers and people understand texts. The task of summarization can aid in this investigation by providing a tool for examining the global structure of a memory representation. This report contains an introduction to the *plot unit* system—a high level representation well-suited for summarizing stories.

Plot units are conceptual structures based on the affective reactions of characters in a story (called *affect states*). By encoding a story into its component affect states and their connections we can derive a *plot unit graph* for the text. Structural features of the graph then reveal which concepts are central to the story, and which are peripheral. Plot units appear to capture the salient aspects of the internal representations used by people when performing tasks requiring a thematuc analysis, and have potential applications in other areas of natural language processing such as inference and natural language generation.

Acknowledgements

Many thanks are due to JoAnn Brooks for her work on plot unit graph analysis and affect state map generation, to John Brolio for his improvemnets to the Plot Unit Graph Generator, and to Kelly Murray for his plot unit graph display program.

Table of Contents

Section 1: Understanding and Representation

1.1 Representational Systems for Text Understanding

If a computer is to be said to understand a story, we must demand of it the same demonstrations of understanding that we require of people. When a person reads a story, an internal representation for that story is constructed in memory. For a computer to read and understand a story, it too must represent the story's content in memory. We can test both human and computer understanding by using various natural language tasks such as answering questions or summarization. Each task will help us examine a different piece of the understanding process and the underlying representation. Question answering provides us with a method for examining the contents of the memory representation, but tells us very little about how it is structured. We can only guess at how the various pieces fit together. Summarization, on the other hand, requires concentration on the central elements of a story while ignoring peripheral information. As such it provides an excellent tool for investigating the global structure of a memory representation.

A variety of techniques for representing the information (both explicit and implicit) in narratives have already been proposed including predicate calculus formalisms (Kintsch, 1974; Woods, 1970), case grammars (Rumelhart & Norman, 1975; Fillmore, 1968; Winograd, 1972; Graesser, 1981) and systems of decomposition into primitives (Schank, 1975; Wilks, 1978; Lehnert, 1979). None of these systems alone can adequately represent the many facets of a large memory structure. Since we often find it necessary to deal with such large representations, there is clearly a need for a new approach. To handle these problems we must turn to multi-layered representations, where the various levels specialize in describing aspects of the text ranging from physical events to general thematic patterns (Dyer, 1983). In this chapter we will present a model for narrative summarization. Our model includes a high-level representational system for story content which is particularly well suited for summarizing stories. This system of *plot units* will allow us to see not only how such a memory representation might be structured, but also how it can be used in a process model for generating summaries.

1.2 Understanding and Inference

What does it mean to understand something we've read? Does it mean the same thing to say a computer "understands" a text? A narrative is more than the sum of its individual sentences; readers supply their own inferences (Reiger, 1975; Schank & Abelson, 1977; Wilensky, 1980), idiosyncratic interpretations and personal belief systems (Carbonell, 1978) during the understanding process. A tremendous amount of knowledge is needed to understand even a simple sentence. For example, to fully understand the sentence "Tod hit Bill." we have to know more than the definitions of the words used. We expect that Bill felt pain and that Tod was probably angry at him. We also know that Tod probably didn't punch Bill in the knee—the upper body is a much more likely target. These *default assumptions* come about as a result of our combined lexical, linguistic and general world knowledge, and may be further refined by idiosyncratic experiences and situation specific information. If we had just read that Tod and Bill were boxers or karate experts, our corresponding assumptions would change dramatically. It has been estimated that the ratio of implicit information derived from a text to explicit information present in the text is something like 8:1 (Graesser, 1981). If a text understander (human or otherwise) is not generating these inferences, we cannot say in what sense that text has been understood.

It is widely conceded that concepts from a text must be stored in memory in some form other than the original sentences, although there is no general agreement as to what this form must be. The necessity of including the inferences generated in story understanding as well as the explicitly present information, has forced us to move away from sentence-driven propositions (Kintsche, 1974) to a more integrated representational scheme (Dyer, 1983). In addition to describing the physical events and situations mentioned in the text, we must also be able to handle inferred goals, likely plans to attain those goals, and affective reactions of characters in the narrative. While sentences will continue to be a necessary starting point, the conceptual information which we must represent will also depend strongly on causal relations and typical character interactions present in the narrative.

1.3 Computer Modelling of Human Understanding

It may at first seem strange to study computers in order to learn more about human understanding and memory representation. Why don't we just study people? The answer to this has many parts. First of all, people are "black boxes." We can ask them to perform some task and observe what they do, but we cannot get inside their heads to see what is actually going on. By constructing computer models of what we believe is happening in a person's mind, we can not only check the model's behavior against the human subject's, but we can also examine its inner workings. Computer models are easy to manipulate. By changing a few parameters we can experiment with the limits of our model and our theory. And we often find that we learn a great deal simply through the process of implementing our model. To program a computer to behave according to our model we must specify each segment of our theory, each rule and each step with great precision. This mandatory level of detail forces us to be very exact about what our model involves.

The cycle of theory formation, implementation and refinement is never ending. After we have implemented our model in a computer and tested our prototypes, we must return again to human subjects for verification. From our observations of human experiments we can then refine our theory and adjust our model. In this way each process helps the other. The knowledge we gain from psychological experiments on human information processing helps us to build better and more capable computer models, and our experiences with these programs give us more understanding of human information processing. In our examination of the summarization process in this chapter, we hope to gain more insight into the phenomena of human and computer understanding.

Section 2: The Plot Unit Representational System

2.1 Affect States

Plot units are constructed from smaller entities called *affect states*. We use affect states to represent a character's mental plans, goals and reactions to external events. Affect states do not attempt to describe subtle or complex emotions; they merely mark gross distinctions between "positive" events (represented by a +), "negative" events (–), and "mental" events of neutral affect (M). At first glance it may seem that such a simplified representation of affect cannot be of much use in narrative processing tasks. This is probably true for tasks like generating inferences or question answering which require in-depth understanding to be successful. But constructing a summary necessitates a loss of detail. While we do not contend that this system is adequate to fully represent the vast range of human emotion, it is nevertheless instructive to see how far this very simple scheme can take us in our search for summarization algorithms.

As we process a story we construct an *affect state map*—a sequence of chronologically ordered affect states for the characters in the story. Each affect state occurs with respect to a single character, so events involving more than one character require multiple affect states. Thus if Jason is in an accident and breaks his arm, we can assume that this event is negative for him, while Linda (who despises Jason) may experience his accident as a positive event.

2.2 Causal Links

We indicate the different relationships between various affect states with pairwise causal links, the second component of plot units and affect state maps. For example, a link which runs from a negative event to a mental state describes motivation, while a link running from a mental state to a positive event describes actualizing a goal.

−
M

link used to indicate motivation

M
+

link used to indicate actualization

To make such distinctions explicit we employ four separate link types: motivation (m), actualization (a), termination (t) and equivalence (e). M-links describe causality behind mental states and a-links indicate intentionalities behind events. T-links indicate a change over time. We use t-links with positive and negative states when the affective impact of an earlier event is displaced by a reaction to a later event. The use of a termination link does not necessarily indicate that the initial event itself has been terminated. It is important to remember that t-links refer to affective reactions and not directly to events. Thus Brenda's marriage to Tom may be a very joyous event for her, but when she discovers Tom is having an affair, her anger may "terminate" her prior happiness, but not the marriage itself. T-links connecting mental states indicate that a prior goal has been displaced, signifying a change of mind. E-links describe multiple reactions to or perspectives on a single event when used with positive and negative states, but represent the reinstantiation of a previous goal when used with mental states.

We impose some syntactic restrictions on the use of causal links to constrain them to our intended meanings. Since we intend m-links to represent the motivation underlying mental goals and plans, they must point to a mental state. Similarly, since a-links indicate intentionalities behind events, they must point from a mental state to a positive or negative event. These two restrictions are the same as those placed by Schank (1975) in his representation of action-state causal chains. In this segment of Conceptual Dependency theory, Schank notes two constraints on how mental states and actions may be linked: 1) states or acts can initiate mental states (equivalent to restrictions on m-links; anything can motivate a mental state) and 2) mental states can be reasons for actions (equivalent to restrictions on a-links; actualizing a goal produces action). T-links and e-links must point from an event to an event, or from a mental state to a mental state as they indicate relationships between like kinds of affect states. These constraints reduce the set of legal pairwise configurations of our three affect states and four link types from 36 possible arrangements to 15 legal arrangements.

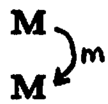

m-link indicates motivation for a goal

(must point to an M-state*)*

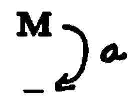

a-*link indicates intentionality*

(must point from an M-state *to* a + *or* −*)*

Causal links have been given a temporal orientation for intuitive convenience. M-links and a-links point forward in time (down the affect state map) from an antecedent to its consequence. With t-links and e-links the pointer goes back in time since each implies a reference to a previous affect state.

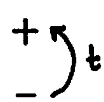

t-link indicates change over time

(must connect like states)

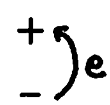

e-link indicates different reactions to a single event

(must connect like states)

2.3 Cross-Character Links

In addition to the four links used to connect affect states for a single character, we also use a cross-character link (c-link) to connect affect states experienced by two different characters. While c-links preserve the temporal order of an affect state map, they have no inherent orientation. Cross-character links can connect any combination of states and events across two characters, and like e-links, their interpretation depends on the specific affect states involved.

Events which result in positive or negative reactions for the second character give us one kind of c-link interpretation. Here we are able work with only two general templates: positive-reaction and negative-reaction.

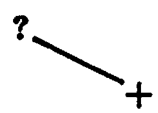

positive-reaction

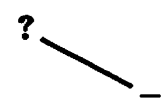

negative-reaction

In labeling such events we do not distinguish among the many different affect states which could occur in the initial position, causing the specified reaction. Different events in the text will yield different configurations, however, each having a slightly different interpretation. A preceding mental state represents a speech act such as a threat or promise, while a preceding positive or negative state indicates a situation involving two characters' reactions to a single event. Thus, the example above of Linda's glee over Jason's broken arm might be represented as:

Often a character will initiate a goal state as a direct response to another character's situation. Configurations of this type make up the second group of cross-character interactions.

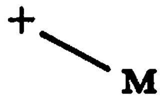

external-enablement

external-problem

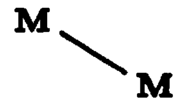

external-motivation

In the case of external-motivation, the resulting mental state usually occurs in response to a request. The responding character may agree to the wishes of the instigator, or may oppose them. This configuration does not commit us to any assumptions about the contents of the two mental states or how they are related. In the cases of external-enablement and external-problem, we have mental states brought on by vicarious events. For example, a desire to celebrate is normally enabled by a positive event while a desire to help out is typically motivated by a negative event.

In the same manner that we represent shared reactions or goals for two characters, we can also indicate how three or more characters respond to a given event. In all such cases we will consider these states to occur in response to a single event. So if three characters share a goal state, we use three M-states c-linked together. If two characters have a positive reaction to some event and a third experiences a negative reaction to the same event, the affect state map would inculde two positive states and one negative state c-linked across. When such a configuration arises, we treat the c-links as if they were "transitive," enabling us to represent the interactions between all pairs of characters using a minimum number of links. So in the example below, this configuration represents not only external motivation between Ken and Sue and between Sue and Debra, but also between Ken and Debra.

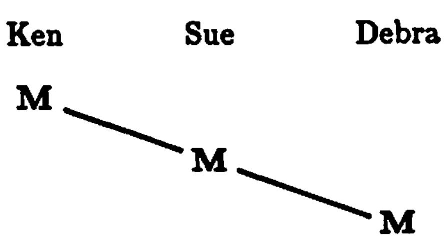

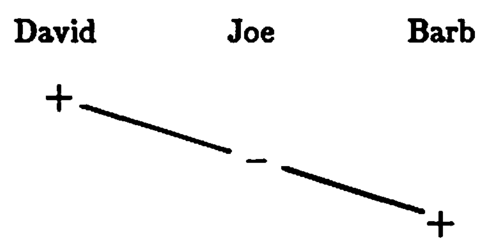

This interpretation becomes important when intervening affects states would otherwise prevent us from realizing a certain relationship was present. Thus in the example above, if David's success angers Joe but thrills Barb, we can recognize Barb's positive reaction to David's triumph only if we permit the transitive interpretation of c-links.

2.4 Primitive Plot Units

The 15 legal configurations of two affect states connected by a single link plus the five two-character pairs listed above form the set of primitive plot units. Each unit's name is meant to be suggestive of its interpretation, although they should not be taken too literally.

See Appendix A for list of primitive plot units

Some Examples

Enablement:	A positive event motivates a goal *You inherit some money and decide to buy a house*
Motivation:	A goal motivates a subgoal *You want to buy a house so you decide to save money*
Success:	A goal is successfully actualized *You want to save money and then you do*
Failure:	A goal is unsuccessfully actualized *You want to save money but end up spending it*
Problem:	A negative event motivates a goal *You get fired and need a job*

Resolution: A positive event terminates a prior negative one

You get fired, but then are offered another job

Loss: A negative event terminates a prior positive one

You are offered a job, but then they give it to someone else

Perseverence: Reinstantiation of a goal

You reapply to a college after being rejected

Change of Mind: A new goal terminates a previous goal

You want to go to Yale, but then decide to go to Harvard

While the set of primitive plot units contains all the possible relations between affect states, they do not, by themselves, give us all the recognition abilities we need. Just as we build the set of primitive plot units from affect states and links, we can now construct more complex plot configurations using the primitive units as building blocks.

2.5 Complex Plot Units

Complex plot units are made up of over-lapping configurations of primitive plot units. But unlike the set of primitive plot units which consisted of every possible arrangement permitted by the syntax of links and affect states, the complex plot units correspond only to situations commonly found in narratives.

Since many stories are about a protagonist's attempts to solve a problem we need to construct plot units to represent the various possible attempts and outcomes. For example, if a character realizes an initial success, only to have it terminated by a later loss we have the *fleeting success* plot unit. If this loss motivates the character to try for the same goal again we have a case of *starting over*. But if an attempt fails, resulting in a change of plans, we get *giving up*.

Fleeting Success

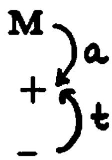

= *success*
& loss

Starting Over

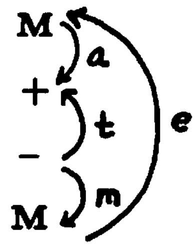

= *success*
& loss
& problem
& perseverance

Giving Up

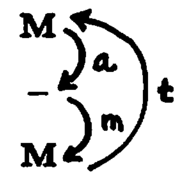

= *failure*
& problem
& change of mind

Other complex plot units dealing with goals and outcomes include *top-level failure* (a subgoal is achieved, but one fails in achieving the top-level goal), both *intentional problem resolution* and *fortuitous problem resolution* (differing only in the intentionality behind the positive outcome) and *half loaf* (a goal is only partially achieved).

Intentional
Problem Resolution

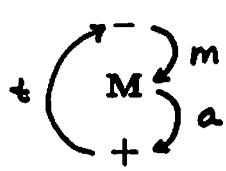

Fortuitous
Problem Resolution

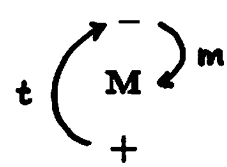

Top-level Failure

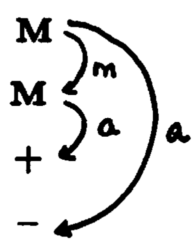

Half Loaf

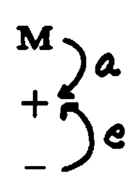

Frequently a character will develop rather intricate plans to achieve a goal. This level of detail can be represented by the *nested subgoals* and *sequential subgoals* plot units.

Sequential Subgoals

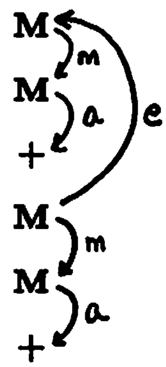

Nested Subgoals

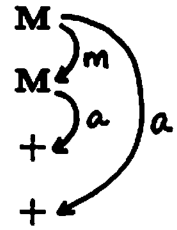

A great many important events occuring in narratives involve multiple characters, so it is not surprising that a large number of complex plot units also involve more than one character. Some of the most common of these are those that involve cooperative agreements and behavior. In the simplest case, a request is made and responded to. The respondent may cooperate *(honored request)*, refuse *(denied request)* or fail in an attempt to fulfill the request *(bungled request)*.

Honored Request

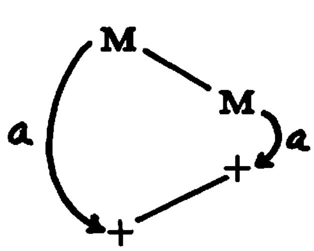

Denied Request

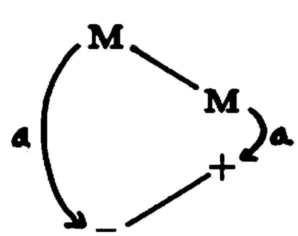

Bungled Request

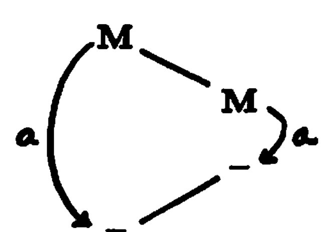

In some cases, the agreement or promise itself is sufficiently interesting to be included in an affect state map. Of course, we also need to be able to represent the eventual outcome of such an agreement—did the respondent fulfill the promise?

Promise

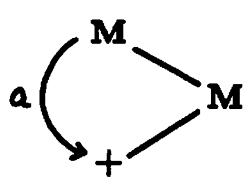

Honored Promise

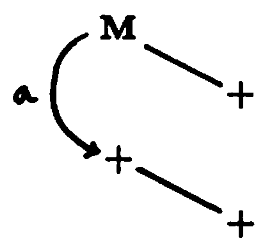

Reneged Promise

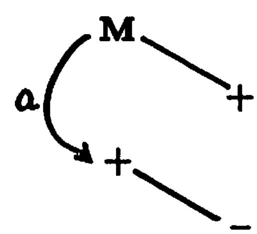

Competition, *retaliation* and *reward* serve as the basis for many storylines and therefore also occur commonly as plot units. In the case of competition, two characters must each be pursuing a unique goal (Obviously, if they both desired the same outcome they would not be in competition with each other.) The critical component of competition is that one person's success must result in the second person's failure. Two drivers attempting to win a race is a clear example of competition. In both the retaliation and reward plot units, an unspecified external event either negatively or positively affects a second character. When this reaction is negative, the responding character instantiates a goal to retaliate. Successfully actualizing this goal produces the desired negative reaction in the first character. Similarly, when the initial event yields a positive reaction, that character forms the goal of rewarding the instigator. This time, achieving that goal results in a positive reaction for the starting character.

Competition

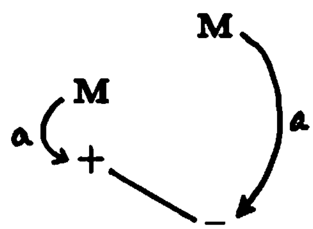

Retaliation

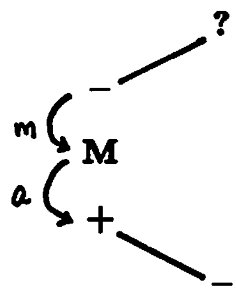

Reward

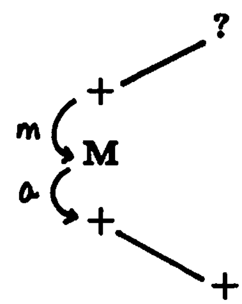

In addition to cooperative and antagonistic reactions, people often respond in unsolicited ways. The act of setting up and carrying out a plan to please or hurt another individual yields either the *kind act* or the *malicious act* plot unit. When the second character acts to alleviate a problem state for the first character, we get *unsolicited help*. In all of these cases we assume that no specific request is made.

Kind Act

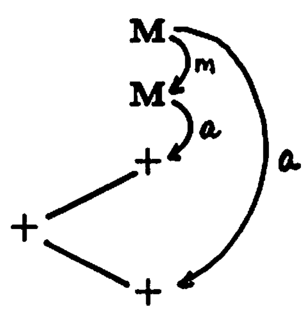

Malicious Act

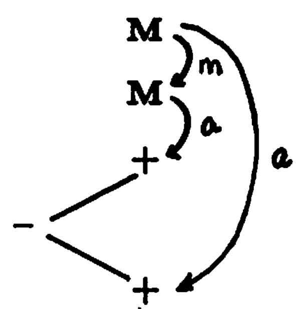

Unsolicited Help

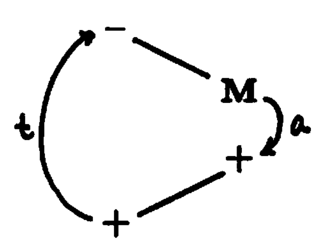

We sometimes come across situations where more than one affect state configuration seems to capture the desired thematic interaction. If one of the drivers mentioned before wins the race only to be disqualified when his opponent files a complaint, we maintain a sense of competition even though the failure component has been replaced by fleeting success. Unsolicited help can also incur a slight modification. If Peter offers to help Bob get his car started because he needs to ride into work with Bob that day, we still have a case of unsolicited help since Bob has not requested any assistance. Here Peter is acting to alleviate both Bob's problem and his own, so we substitute intentional problem resolution for success in our definition of unsolicted help to arrive at the new version of this plot unit. These "flexibly defined" plot units capture the thematic sense of an interaction even when slight variations occur in the action. The substitutions used here—fleeting success for failure and intentional problem resolution for success—are common ones and can be used in many different situations, while still preserving the gist of a particular plot unit.

Competition-2

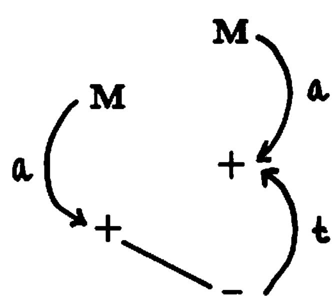

Unsolicited Help-2

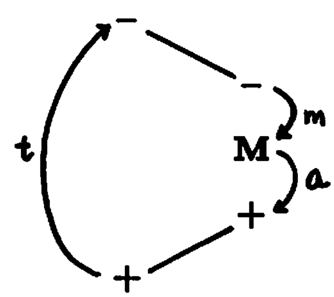

2.6 An Example

The best way to see how the different components (affect states, causal links, primitive and complex plot units) are used is to walk through a quick example. By comparing the story, the affect state map and the explanation below, one should be able to get a feel for the system. (For a more detailed treatment of guidelines and heuristics used in constructing affect state maps see Brooks, 1984.)

The History Project

Jennifer and Allison both had a crush on Mark. Since they were all in the same history class, it was natural for Mark to ask Jennifer to help him with his class project. Jennifer was afraid, however, that if Mark thought of her only as a study partner, he wouldn't consider her for a date, so she turned him down. Mark didn't know what to do. With only three days left before the project was due he was afraid he'd never get it done in time. Allison, overhearing the conversation with Jennifer, volunteered to help Mark with his project. Naturally, Mark was quite relieved. As they worked together on the project, Mark discovered that he really liked Allison's company, so he asked her to go with him to the big rock concert that weekend. When Jennifer found out, she realized how mistaken she'd been.

Affect State Map

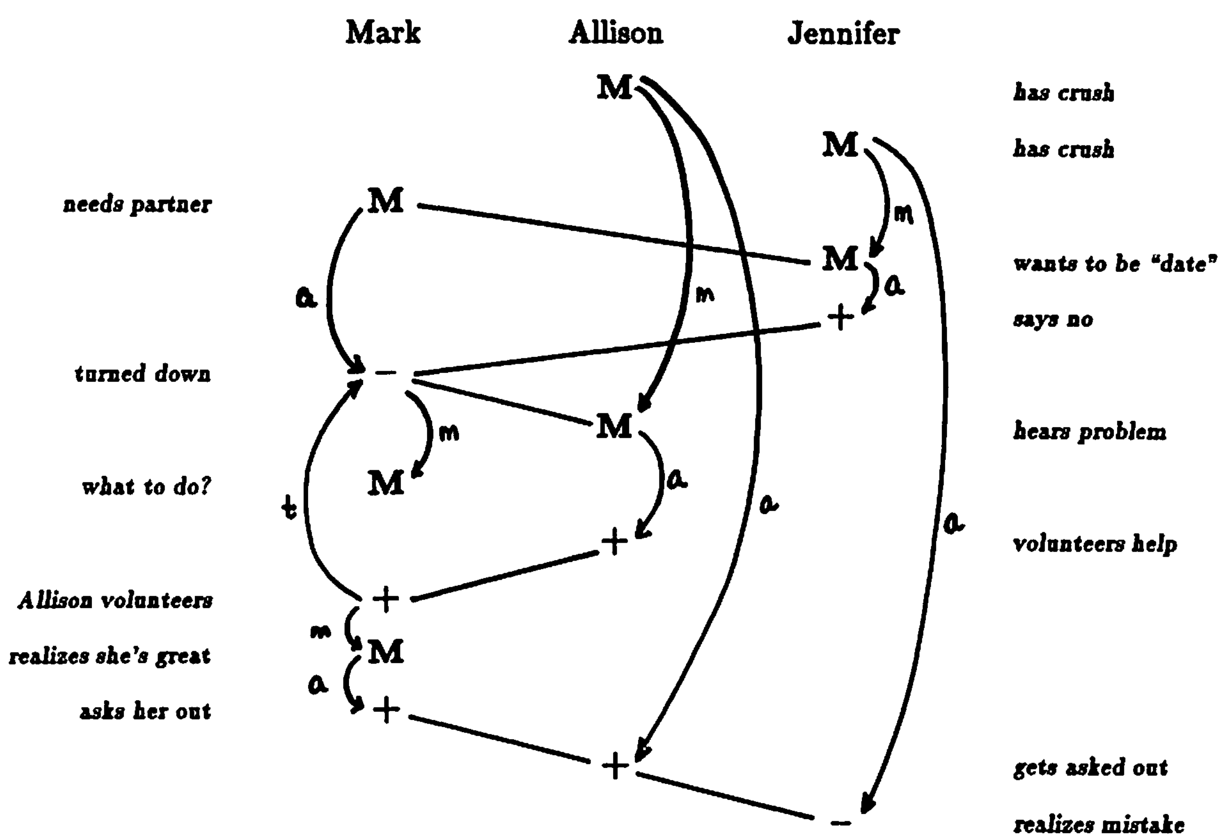

Explanation

The story begins by stating a goal for each of the three characters. Jennifer and Allison want to date Mark while he wants to complete his history project. We represent these goals with mental states. When Jennifer refuses to help Mark we have an example of a denied request, motivated in this case by her previously stated goal. This instantiates a problem for Mark since he now has no one to work with. Allison's crush on Mark motivates her offer of assistance. Since Mark did not plan this as a solution or request her help we represent this with the unsolicited help and fortuitous problem resolution plot units. Mark's gratitude prompts him to reward Allison by asking her out. With this act Allison's initial goal of dating Mark is achieved, thwarting Jennifer's plans and establishing her as the victor in their competition over Mark. We further note that although each girl realized her subgoal, Allison was successful in achieving her top-level goal (giving us the nested subgoals unit) while Jennifer was not (resulting in top-level failure.)

Section 3: Summarization from Plot Units

Now that we have some familiarity with the plot unit representational system, let us examine how such a scheme can be used for narrative summarization. This process involves several steps beyond constructing an affect state map. These are (1) recognizing all the plot units present in the map, (2) sifting out the "top-level" plot units, (3) building a connectivity graph and (4) using structural features of this graph to identify the components from which we can generate a summary. In this section, we will present each of these steps in more detail as well as discuss the various factors that can influence the structure of this graph, and therefore the summary itself.

3.1 Identification of Plot Units

The recognition of plot units proceeds from the top to the bottom of the affect state map (Lehnert & Loiselle, 1983). As we encounter each affect state in the map, we form predictions for the states and links we expect to encounter next. These predictions take the form of demons which attempt to recognize plot units by matching the states and links in the map with templates for the anticipated plot units. We form predictions only when we have seen enough structure to justify such expectations. Thus when we encounter an M-state in the map, we only form predictions for the primitive plot units with an initial M-state. When we recognize a primitive unit such as success, we form predictions for complex plot units which begin with the success configuration. By ordering our predictions hierarchically, we minimize the search time spent considering possible structures. Once we have identified all the plot units in a given affect state map we proceed to construct the *plot unit graph*.[1]

3.2 The Plot Unit Graph

As we have seen in the "History Project" story in section two, plot units can overlap with one another at shared affect states. When a plot unit totally envelops a smaller unit

[1] In actuality, the identification of top-level plot units (discussed in the next subsection) occurs concurrently with the recognition of all the units present in the affect state map.

we say that the smaller unit is *subsumed* by the larger plot unit. (We have seen this before—complex plot units subsume the primitive units from which they are built.) We use the "top-level" plot units (those units not subsumed by any other unit) to construct the *plot unit graph*.[2] Each top-level plot unit is represented by a node in the graph. Nodes are linked together when two plot units overlap at at least one affect state. We say that two such nodes are *related*. Two plot units or nodes are *connected* if they are related, or if there is some path through any number of related units which joins them. This partial affect state map from the "History Project" story shows how plot units can overlap and subsume one another.

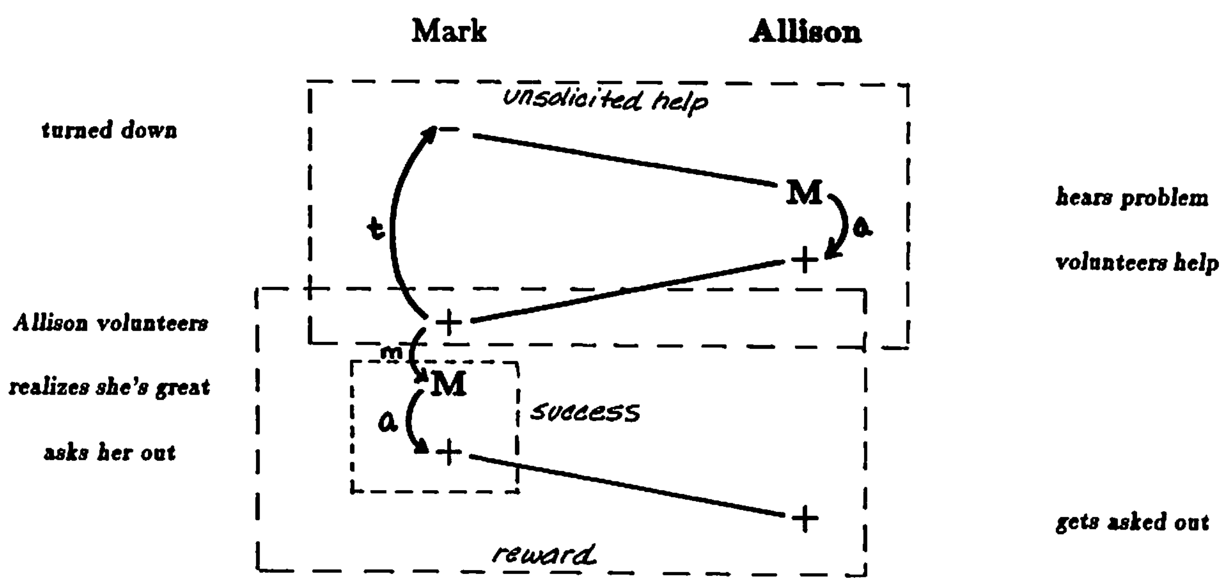

In this map the reward plot unit subsumes the smaller success unit while unsolicited help and reward overlap, sharing two positive affect states. The plot unit graph for this partial map contains two nodes linked together, representing the two overlapping units. Success does not appear on the graph since it is subsumed by the reward plot unit.

By using structural features of the resulting graph we are able to identify the critical

[2] Primitive plot units may still rise to top level in a given affect state map if they are not subsumed by a larger unit, although by convention, the dyadic (cross-character) primitives are not included in the plot unit graph, even when they appear as top-level units.

nodes in the plot unit graph. It is these nodes which contain the units central to the narrative and from which we can generate a summary. The process of identifying these critical nodes forms a retrieval algorithm for summarizing the given narrative. Different types of stories will produce different plot unit graph structures. We have identified three core classes, each with corresponding retrieval algorithms (Lehnert, 1983).

3.3 Structural Classes of Plot Unit Graphs

One class of graphs exhibit unique nodes of maximal degree. We label such a graph a "simple cluster." While this class seems to be restricted to smaller graphs, we can reliably look to such pivotal nodes for the concepts most central to the story as a whole. Here we use the maximal node to generate a baseline summary, augmented by the concepts contained in the immediate relatives of that node. Let us look at a simple example to see how this might proceed.

The Broken Engagement

> Doug was thrilled when Becky accepted his engagement ring. But when he found out about her father's illegal mail-order business, he felt torn between his love for Becky and his responsibility as a policeman. When Doug finally arrested the old man, Becky called off the engagement.

Affect State Map

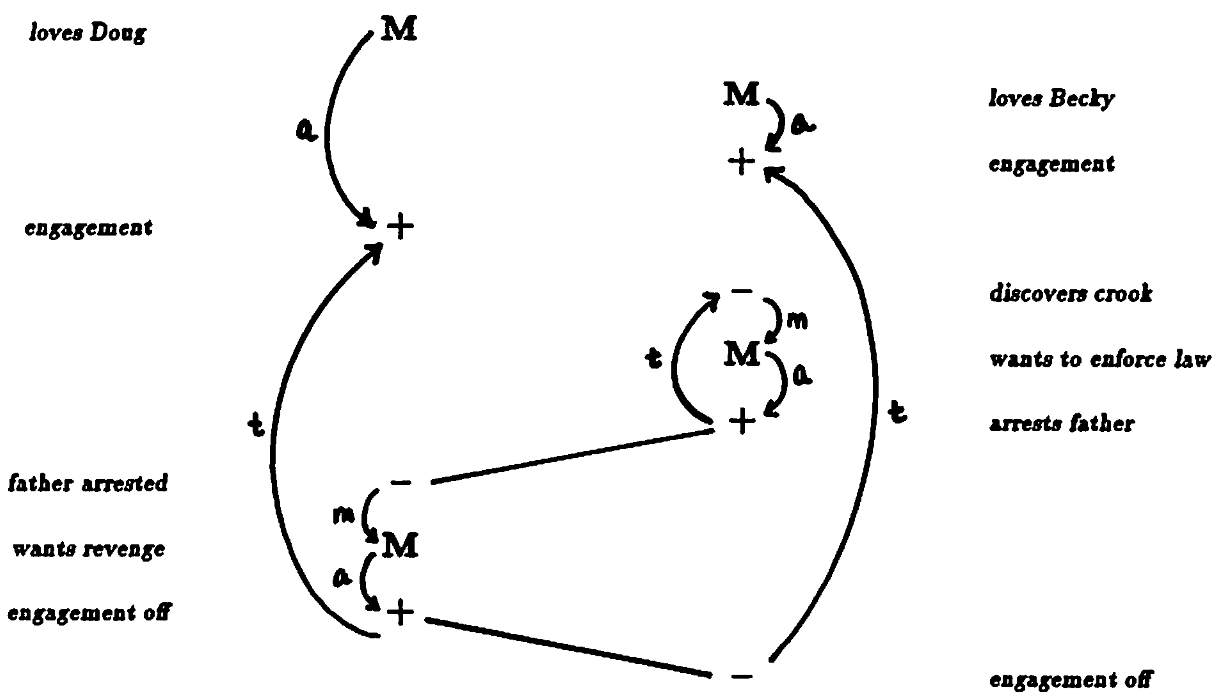

Plot Unit Graph

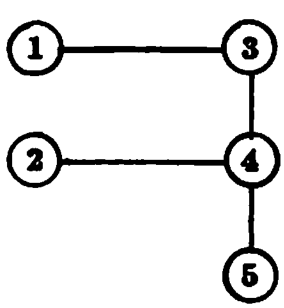

1. Success (Becky)
2. Intentional Problem Resolution (Doug)
3. Positive Trade-Off (Becky)
4. Retaliation (Becky against Doug)
5. Fleeting Success (Doug)

The plot unit graph for this story is a simple cluster. The node representing the retaliation plot unit has maximal degree and is therefore the central concept for our summary. This seems to match our intuitive sense of the story's main point as well, but a summary containing only that idea would be too weak. We need to augment this unit with its immediate neighbors to produce an acceptable summary. Leaving out any one of these units out will weaken the summary:

Becky got back at Doug for hurting her.

(retaliation only)

When Doug arrested Becky's father, she interfered with his wedding.

(no trade-off for Becky)

When Doug arrested an old crook, Becky called off their engagement.

(no retaliation)

When Becky's father was arrested, she called off their engagement.

(no fleeting success for Doug)

But a summary that includes all four plot units provides an accurate description of the story:

When Doug arrested Becky's father, she called off their engagement.

(all units present)

As stories get longer and more complex, the top-level plot unit graphs follow suit. Here we often encounter graphs with multiple pivots (nodes of maximal degree). This class of graphs can be further divided into smaller subsets, each group still having its own summarization algorithm.

In one subset are the graphs where the maximal nodes provide the key concepts for summarization, much like the simple cluster graphs discussed above. This is especially common when the pivotal nodes are adjacent to one another. The "History Project" story falls into this class.

The History Project

> Jennifer and Allison both had a crush on Mark. Since they were all in the same history class, it was natural for Mark to ask Jennifer to help him with his class project. Jennifer was afraid, however, that if Mark thought of her only as a study partner, he wouldn't consider her for a date, so she turned him down. Mark didn't know what to do. With only three days left before the project was due he was afraid he'd never get it done in time. Allison, overhearing the conversation with Jennifer, volunteered to help Mark with his project. Naturally, Mark was quite relieved. As they worked together on the project, Mark discovered that he really liked Allison's company, so he asked her to go with him to the big rock concert that weekend. When Jennifer found out, she realized how mistaken she'd been.

Using the affect state map in section 2.6, we can derive the following plot unit graph:

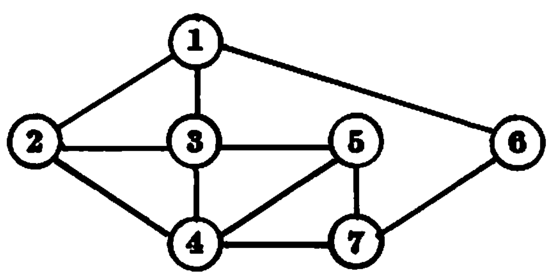

1. Denied Request
2. Fortuitous Problem Resolution
3. Unsolicited help
4. Reward
5. Nested Subgoals
6. Top-Level Failure
7. Competition

Nodes 3 and 4 (unsolicited help and reward) both have maximal degree and therefore are the critical nodes for this graph. A simple summary that contained only these units might be:

> Mark asked Allison out (*reward*) after she volunteered to help him with his history project (*unsolicited help*).

This summary feels better than the one generated from the critical node in the "Broken Engagement" story since we are starting with two key nodes this time. If we want to, though, we can fill this out with these nodes' relatives (denied request, fortuitous problem resolution, nested subgoals and competition). A summary including all these units might be:

> Allison beat out Jennifer (*competition*) and got her wish (*nested subgoals*) when Mark asked her out (*reward*) after she volunteered to help him with his history project (*unsolicited help*). This got him out of a jam (*fortuitous problem resolution*) caused by Jennifer's earlier refusal to work with him (*denied request*).

This graph is a good example of how we can use connectivity to select the nodes from which we will generate a summary. Including all four relatives of our two critical nodes forces us to include six of our seven top-level units in the summary. We note that although each of these relatives has degree three, only two (fortuitous problem resolution and nested subgoals) are linked to both of our critical nodes. This suggests a third possible variation for this graph using only the two critical nodes and their most tightly connected relatives:

> Allison got her wish (*nested subgoals*) when Mark asked her out (*reward*) after she got him out of a jam (*fortuitous problem resolution*) by volunteering to help him with his history project (*unsolicited help*).

In other cases we see graphs where the pivotal plot units partition the graph into two subgraphs. Here we find that the nodes on the boundary between these units are the critical ones. We can use the degree of these boundary nodes as a general guide in selecting the essential units when many such boundary nodes exist.

This unusual structure shows up nicely in O. Henry's *Gift of the Magi*. This is a story about a young couple who want to buy each other Christmas presents. They are both very poor. Della has long, beautiful hair, and Jim has a prized pocket watch. To get money for the presents, Della sells her hair and Jim sells his watch. Then, she buys him a gold chain for his watch, and he buys her an expensive ornament for her hair. When they find out what they've done, they are consoled by the love behind each other's sacrifices.

The story's affect state encoding exhibits an extreme symmetry:

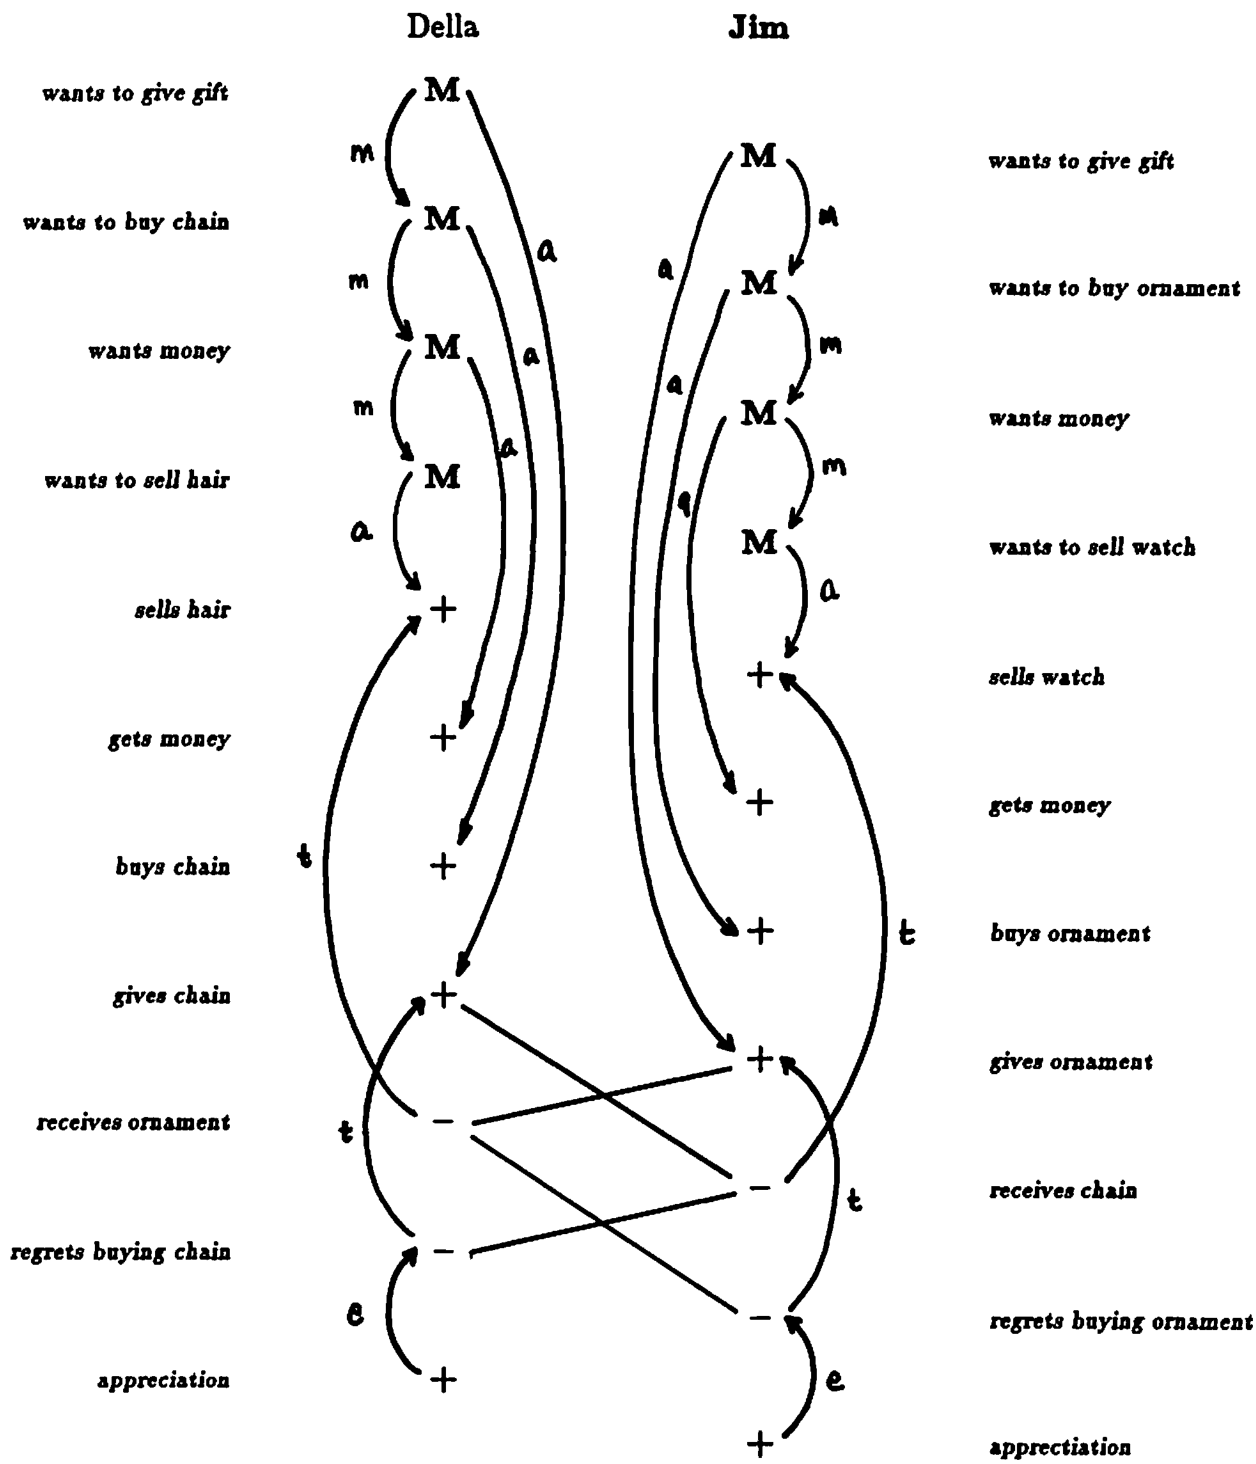

This symmetry is quite naturally reflected in the plot unit graph:

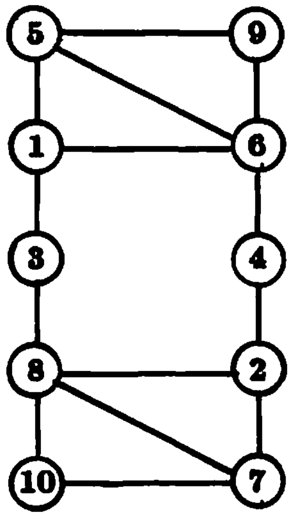

1. Nested Subgoals (Della)
2. Nested Subgoals (Jim)
3. Fleeting Success (Della—receiving gift)
4. Fleeting Success (Jim—receiving gift)
5. Fleeting Success (Della—giving gift)
6. Regrettable Mistake (Della to Jim)
7. Fleeting Success (Jim—giving gift)
8. Regrettable Mistake (Jim to Della)
9. Hidden Blessing (Della)
10. Hidden Blessing (Jim)

Because of this symmetry, we see the same units appearing for both Della and Jim. There are two nodes of maximal degree in the graph (regrettable mistake for both characters). These nodes and their immediate families form two distinct subgraphs with nodes 1, 2, 3, and 4 on the boundary. The two nested subgoals units represent each character's plans to sell a precious possession in order to buy their spouse a gift. The two fleeting success units on the boundary are those for Della and Jim each receiving a gift they know they cannot use. Although these are certainly key events in the story, they are not sufficient for a good summary. Because of the story's complexity in this case, we might want to use all the top-level plot units in a summary. We can use our retrieval algorithm to order their appearance in the summary, placing our critical nodes early in the story:

> A woman sold her long locks of hair so she could buy her husband a watch chain for Christmas (*nested subgoals (1)*). But when she gave him the chain she found out that he had sold his watch (*fleeting success (4 & 5)*) so he could but her a comb for her hair (*nested subgoals (2), fleeting success (3 & 7)*). Initially they regretted their expensive gifts (*regrettable mistake (6 & 8)*), but then they realized how much love was signified in the sacrifices made (*hidden blessing (9 & 10)*).

Our third class of graphs is made up of very large graphs (more than 50 nodes) composed of subgraphs that can be separated by deleting a single node called an *articulation point*. Any such articulation point whose removal would result in the separation of at least ten percent of the graph, becomes a pivotal unit for the story. When multiple articulation points exist, the path connecting all such nodes becomes the basis for a summary. The shortest such path is usually preferred, although the degree of the nodes along the path can also be a factor. (See Lehnert, Alker & Schneider (1983) for a detailed discussion of one such graph.)

3.4 Influences on Plot Unit Graph Structure

Several modifications in how an affect state map is constructed can change the plot unit graph, and possibly the resulting summary as well. Such alternate encodings result from two main variations: (1) level of detail and (2) inferences made by the story understander (Brooks, 1984).

Obviously, if we include more detail in an affect state map we will produce a larger map, and also a larger plot unit graph. Our main concern is whether such a change will affect the critical nodes of the graph, thereby changing the summary. To answer this we must examine how we can vary the level of detail in an affect state map. We will use the story of the "Czar's Daughters" as our example.

The Czar's Daughters

> Once there was a Czar who had three lovely daughters. One day the three daughters went walking in the woods. They were enjoying themselves so much that they forgot the time and stayed too long. A dragon kidnapped the three daughters. As they were being dragged off they cried for help. Three heros came and fought the dragon and rescued the maidens. Then the heros returned the daughters to their palace. When the Czar heard of the rescue, he rewarded the heros.

First we will present an affect state encoding corresponding to the level of detail we have been using up to now:

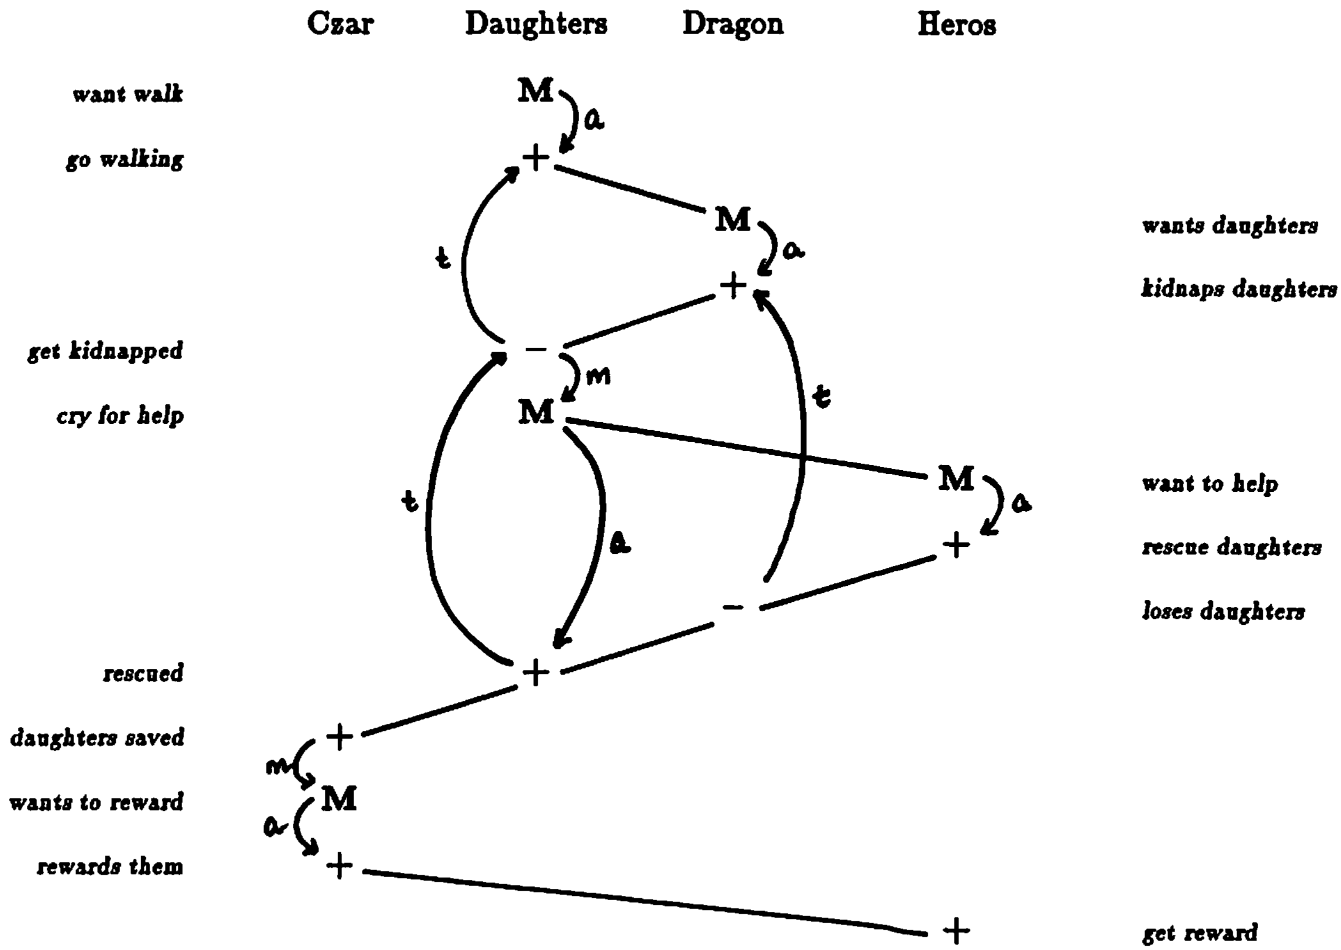

The resulting top level plot unit graph is:

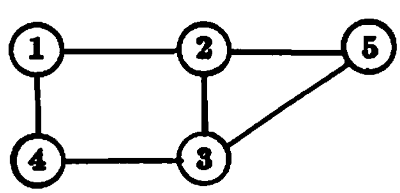

1. Competition (Daughters & Dragon)
2. Competition (Heros & Dragon)
3. Honored-Request (Daughters & Heros)
4. Intentional Problem Resolution (Daughters)
5. Reward (Czar & Heros)

We have two adjacent nodes of maximal degree (competition between the heros and the dragon and honored request between the daughters and the heros). Following our summarization algorithm for this case we would select these as our critical nodes on which to base a summary. Our main interest in this section will be to see if modifications to the affect state map will change our choice of critical nodes and thus the summary itself.

We have followed our usual convention of representing intentional acts with the success unit. The girls' going for a walk, the dragon's capturing the girls, and the heros' conquering of the dragon are all represented in this way. Some of these events can easily be broken down into several acts increasing the level of detail included in the map. For example, we know that the daughters went for a walk and then stayed out too late. We can represent this as:

goal:	want to go for a walk	M
event:	go walking	+
goal:	want to stay longer	M
event:	stay longer	+

chaining the events together. Similarly, we could expand upon the heros' battle with the dragon. First they hear the daughters' cry for help. Deciding to help them motivates a goal of defeating the dragon, which they then carry out.

goal:	want to save daughters	M
goal:	want to defeat dragon	M
event:	defeat dragon	+
event:	daughters are saved	+

When we make these two changes in the affect state map, we end up with the following plot unit graph:

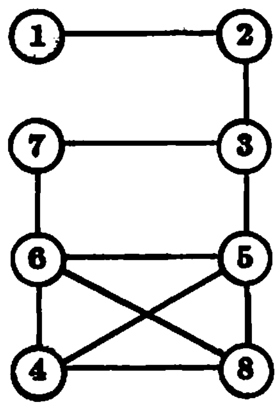

1. Success (Daughters)
2. Enabled Success (Daughters)
3. Competition (Daughters & Dragon)
4. Nested Subgoals (Heros)
5. Competition (Heros & Dragon)
6. Honored Request (Daughters & Heros)
7. Intentional Problem Resolution (Daughters)
8. Reward (Czar & Heros)

We have added a few new units to the plot unit graph, but we notice that honored request and competition are still our pivotal units. The resulting summary would probably not differ markedly from one generated from the original plot unit graph.

The other main factor to consider involves the inferences made by the reader or story understander. The three t-links included in the previous graph reflect one type of inference made by the reader. In this encoding, we infer that being kidnapped terminates the daughters' joy in walking, being defeated terminates the dragon's glee over capturing the daughters, and being rescued terminates the daughters' unhappiness from being kidnapped. These are all reasonable assumptions, yet a story understander may not make them all. When we leave them all out we obtain the following top-level plot unit graph:

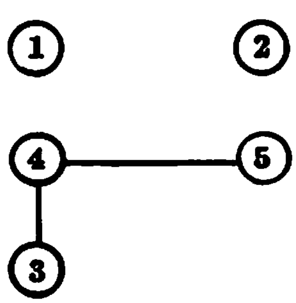

1. Success (Daughters)
2. Success (Dragon)
3. Success Born of Adversity (Daughters)
4. Honored Request (Daughters & Heros)
5. Reward (Czar & Heros)

In this version we have lost many of the complex plot units (here replaced by the two primitive success units) and much of the connectivity as well, although we still retain the honored request as a pivotal node. Deleting the three t-links provides a good example of how drastically a graph can change in appearance as the result of including or excluding certain inferences. But we also can note here the robustness of this representation. Even with such a major change in the top-level plot unit graph, we still select the same concepts as being crucial to the narrative.

Often inference and level of detail are closely intertwined in an affect state map. One reader may infer unspecified motives behind events whereas another may only include the affective impact of the events alone. The character of the dragon is a good subject to experiment with for this modification. At the two extremes we can infer either no motivation behind the dragon's acts, or conclude that each event was motivated by some mental goal. In the first case we might hypothesize that kidnapping fair maidens is a purely instinctive act for dragons, as is fighting any young heros who try to rescue them. In the second case we infer that the dragon had instantiated goals both to first capture, and second, retain the daughters.

The affect state representation for the dragon's character would then correspond to one of these maps:

No Motivation

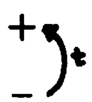

captures daughters

loses daughters

Extended Motivation

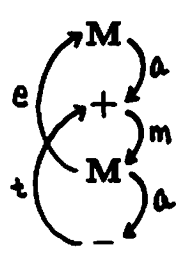

wants daughters

captures daughters

wants to keep daughters

loses daughters

Note that we consider the two inferred goal states to be enough alike to consider one a reinstantiation of the first, thus setting up a perseverance plot unit. To see how these changes might affect our summary, let's examine the corresponding plot unit graphs obtained from each new version:

No Motivation

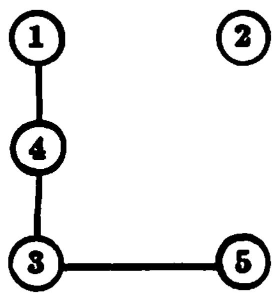

1. Fleeting Success (Daughters)
2. Loss (Dragon)
3. Honored Request (Daughters & Heros)
4. Intentional Problem Resolution (Daughters)
5. Reward (Czar & Heros)

Extended Motivation

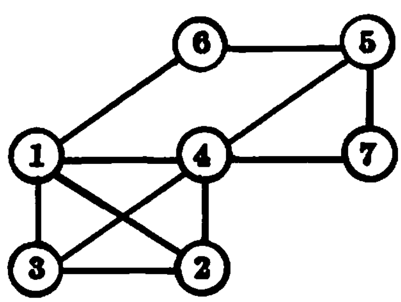

1. Competition (Daughters & Dragon)
2. Perseverance (Dragon)
3. Enablement (Dragon)
4. Competition (Heros & Dragon)
5. Honored Request (Daughters & Heros)
6. Intentional Problem Resolution (Daughters)
7. Reward (Czar & Heros)

With no motivation, we lose both competition units. With both goal states inferred we add several primitive units for the dragon. In the first case, the honored request unit remains a pivotal node, although it now shares that spot with the daughters' intentional problem resolution. In the second case, the competition units return and once again take a central position. Although honored request still has a fairly high degree, it is not as critical to this graph as the two competition units.

3.5 Summarization Research

In this section we have shown how to identify the key concepts in a narrative using structural features of the plot unit graph. The production of natural language summaries from the information contained in the pivotal units is a very interesting but distinct problem which we will discuss in the final section. Identifying appropriate retrieval algorithms and encoding heuristics are some of our major concerns in our current research effort.

To date we have examined approximately 25 stories in detail. For each text we have constructed many different affect state maps, each corresponding to the various inferences or levels of detail that could be used by different readers. The plot unit graphs generated for these stories show very favorable results. Most of these graphs fit nicely into one of the structural classes described above, with the corresponding retrieval algorithm successfully identifying the key concepts for summarization. Our results also support the robustness of the representation. despite the wide variety of assumptions used in their encoding, the identification of critical node has remained quite stable within each story.

Future research plans call for us to expand our library of stories and plot unit graphs, and to continue our study of how encoding variations may influence summarization. We also plan on refining our classification of plot unit graph structures. We believe that such a taxonomy can serve as a basis for establishing equivalence classes for narratives based on summarization algorithms (Lehnert, 1983).

Section 4: Evidence for the Psychological Validity of Plot Units

In addition to the task of summarization, the plot unit memory representation lends itself to many other narrative processing tasks. In this section we will see how affect analysis in general, and plot units in particular, can be used in computational and cognitive models of comprehension, inferencing, memory retrieval, reminding, and story generation and classification.

4.1 Inference Generation

The plot unit representation system grew out of experience with the BORIS story understanding system (Lehnert, Dyer, Johnson, Yang and Harley, 1983) and its analysis of affect. After first examining Roseman's (1979) model for representing affective states, it was determined that higher level knowledge structures were needed to handle the necessary inferences (Lehnert, 1981). Initial efforts in this direction led to the development of Thematic Affect Units, or TAU's, (Dyer, 1981), but plot units also contain a great deal of information which can aid inferencing.

Much of what we know about social and goal relationships is thematic in nature, and therefore relatively invariant across different situations. For this reason, such information is quite useful in text understanding. Plot units capture knowledge about social and goal relationships which is not dependent on the particular activity or situation involved. Recognition of a particular plot unit allows us to infer this additional information, aiding the understanding of events already processed as well as creating expectations for future goals, actions and emotional reactions of the characters. Such inferences will almost always be valid, regardless of the specifics of that interaction.

For example, once we recognize a situation to be a competition, we have a great deal more information available to aid in understanding the text. We can predict that the loser will probably feel angry or disappointed. It is also likely that the two opponents are not friendly towards one another (or if they were, they may not be anymore). Since such inferences are also thematic, they will need to be integrated with situation specific knowledge. We would expect a retaliatory strike from the loser to take a different form when two men fight over a woman than when two children fight over a candy bar, although such revenge would not come as a surprise in either case. In order to achieve a complete understanding of the story, we need to include inferences from both levels (Dyer, 1983).

4.2 Memory Retrieval

Remembering a story that contains more than a few sentences requires condensing information and selectively ignoring details. The plot unit representation accomplishes both of these tasks. If such a knowledge structure is indeed used in understanding and remembering narratives, we would expect its influences to be felt during retrieval as well. In fact, the connectivity of plot units has been found to be a better predictor of retrievability than position in a structural hierarchy (Lehnert, Black & Reiser, 1981).

In this experiment, subjects were asked to produce a short written summary for one of three variations of a particular narrative. Each variation was encoded using both the plot unit (causal connectivity) and the story grammar (position in a structural hierarchy) systems. The experimenters identified predictions made by both models regarding the retrievability of the concepts in memory and compared those predictions to the data. Plot units were found to be a better predictor of a concept's inclusion in a summary than were story grammars.

4.3 Reminding and Classification

Closely related to retrieval from memory is the process of reminding. What sort of memory structure allows us to be reminded of other experiences? Schank has explored (1979, 1982) how an experience might remind one of another experience, and how a story might remind one of another story, suggesting that the thematic structure of the story or experience could play an important role. When a thematic pattern is recognized, other stories or experiences processed using that thematic structure may be brought to mind. In text understanding, reminding can aid in making predictions about what is likely to occur. In this respect reminding is like an inference process. Unlike inferencing, however, we must extrapolate from individual past experiences to the current situation. For example, suppose we had processed several stories involving retaliation after a double-cross. If we are then asked to read another story with a double-cross, we should expect to be reminded of the previous stories, and would probably form a prediction for a retaliation in this case, also.

We can be reminded of previously read stories in many different ways. Two such possible bases for reminding are story content and story structure. This phenomenon allow us to classify stories accordingly. We can group stories according to the kinds of events which occur (content). For example, we can divide stories into those with successful and unsuccessful exection of a plan. A story's structure may also be used for comparison. Stories with deeply nested goal chains might be distinguished from those with a simple sequence of events. At a higher level, stories with a single critical happening could be separated out from those with multiple key events. Plot units may be useful in classifying narratives on both contentive and structural grounds.

In a clustering experiment, Reiser, Lehnert and Black (1981) asked subjects to sort 36 stories into groups with the "same kind of plot." Many dimensions could conceivably influence the subjects' judgement in grouping the stories (type of plan generated by the protagonist, contextual settings, desirability of the story situation, etc.), yet the six clusters of stories found in the data corresponded very well with the six groups of stories predicted by the plot unit representations. Further, some of the subjects' labels accurately reflected the gist of a plot unit (e.g. "broken promise" and "revenge" for the plot units reneged-promise and retaliation, respectively) at the same level of abstraction. Higher level clusters corresponded to thematic judgements of a more abstract nature than that represented by plot units. At this level subjects made discriminations based on factors such as the nature of the outcome and the "fairness" of the protagonist.

In another series of experiments, Gee and Grosjean (1984) used spontaneous pausing during the telling of a story to determine if such empirical data could reflect narrative structure. Pausing during sentence breaks was found to be highly correlated with the importance of these breaks as predicted by a plot unit analysis of the story. Gee and

Grosjean caution us that any model we now have of narrative structure can be only a rough approximation, but hope that pausing data may help refine these theories. We believe that by exploring the different strategies used to summarize various stories, we can develop a taxonomy of narrative complexity where stories whose plot unit graphs can be summarized by the same retrieval algorithm would form an equivalence class.

Higher level structures such as TAU's (Dyer, 1981) and MOP's (Schank, 1979) may also serve as a basis for reminding. These structures involve more complex thematic patterns than those represented by plot units and may be more useful for encoding the major theme of a story. Dyer points out that these structures are often nicely captured in adages. A common thematic pattern involving a planning failure is "executing a plan when it is too late to do any good." The adage for this would of course be "closing the barn door after the horse has fled."

4.4 Story Generation

Plot units may also serve as guides for writing stories. The use of high-level thematic structures such as plot units and TAUs as the basis for a story writing system has not yet been explored, although both representations seem to offer reasonable starting points for story generation. Recent theories of writing suggested that the writing process consists of the stages of idea generation, translation and editing (Hayes & Flower, 1980; Bruce, Collins, Rubin & Gentner, 1980). Hayes and Flower (1980) found that the idea generation phase consists of the production of conceptually related chunks of ideas. Plot units may provide the right kind of "chunks" for such idea generation. For example, if one decides to write a story about a custody fight involving the retaliation of the loser, we would expect to see the various components of the retaliation unit (the triggering incident, the plan, and the act of revenge) present in the story.

In a related experiment, subjects were asked to write stories thematically similar to a set of prototypes (Reiser, Black and Lehnert, 1983). Although the subjects' stories varied widely in setting and context, when analyzed into their component plot units they were largely composed of the same units as the prototype stories. Subjects often embedded a

plot unit from the prototype story into a larger and more complex unit in their ve
For example, the threat plot unit was often incorporated into an ineffective coercion .
When the prototype narratives contained a pivotal plot unit, that particular unit appeai very frequently in the subjects' versions. Subjects also tended to focus more on the protagonist's plans than on the motivating situation, so the plot units corresponding to the protagonist's actions appeared more often in the subjects' stories, than those relating to the establishment of the problem situation itself.

4.5 Conclusion

Plot units seem to well represent the memory structures used in understanding, remembering, summarizing, composing and relating narratives. Plot unit representations appear to capture the salient aspects of whatever internal representation humans use when asked to perform some task requiring a thematic analysis. People use these kinds of concepts when dealing with narratives, whether writing original stories or remembering previously read texts.

Plot units also provide a useful level of representation for computational models of human memory and natural language processing. Thematic structures such as TAU's, MOP's and plot units have been used in a wide variety of computer models ranging from a program which models long-term reconstructive memory (Koldner, 1983) to one which handles the in-depth understanding of stories in a "soap opera" domain (Lehnert, Dyer, Johnson, Yang & Harley, 1983). Plot units in particular were developed with the goal of achieving a process model of narrative summarization and have been implemented in such a program (Lehnert & Loiselle, 1983). Thematic structures such as these will certainly continue to play an important role in such computational models. The evidence presented above suggests that such a thematic analysis permeates the processing of narratives, regardless of task.

Section 5: An Overall View

5.1 Where do Plot Units Fit in?

The plot unit system resides on one level of a multi-layered representational scheme. Since they use a thematic level of description, plot units work well for examining larger, structural aspects of a narrative, but are ill-suited for analyzing the fine details of a story's content. Plot units can, however, play an important role in an integrated representational system. Thorough comprehension of a narrative involves understanding at all levels.

Because affect states are based on information about plans, goals and themes, affect analysis will not be appropriate for expository text or passages. Using this approach we will not be able to handle descriptions of sunsets or comparisons of different forms of government. In fact, it is not clear that people can comfortably summarize such texts, either. The applicability of plot units to a given text might provide us with a means of classifying types of texts, or at least in distinguishing expository text from narratives.

5.2 The Plot Unit Graph Generator

We have implemented the plot unit model for summarizing stories in a computer program. This program—the Plot Unit Graph Generator (PUGG) (Lehnert & Loiselle, 1983)—handles the conceptual portion of narrative summarization. Its level of representation (plot units) must be built on top of lower levels of memory representation. Much work has been done developing these lower levels of memory representation, in particular the BORIS system (Lehnert, Dyer, Johnson, Yang & Harley, 1983) provides a good picture of what such representations must entail. We will not discuss such systems here except to note that the knowledge structures they employ (see, for example, Schank & Abelson, 1977) provide us with enough information for constructing affect state maps. PUGG accepts input in the form of an affect state map, as opposed to English prose.

Let's look at an example to see how PUGG processes a story. Recall that an affect state map is a chronological ordering of the simple affective reactions for the main characters in a story. By comparing the input to and output from PUGG, we can get a good feel for what it accomplishes. Consider the following story and affect state map:

> John wanted to make a lot of money so he decided to go to college and become a doctor. While in college he decided that he liked working with computers better than he liked working with people so he changed his mind and decided to become a computer professional. When he graduated he looked for a position with a high-ranking computer firm and landed a good job at IBM. John ended up making a doctor's salary anyway.

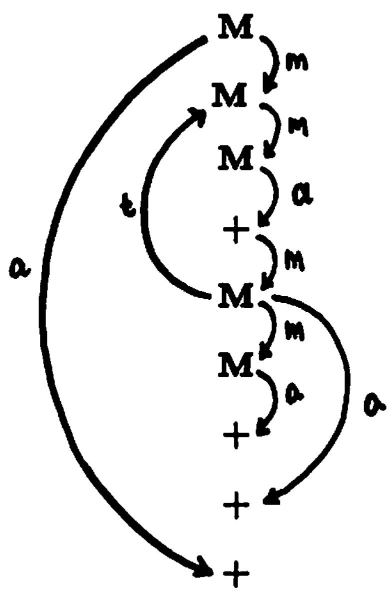

wants to make a lot of money

wants to be a doctor

wants to go to college

goes to college

wants to be a computer professional

want to work with a computer firm

gets a job with IBM

becomes a computer professional

makes a lot of money

We can see that we start out with a simple list of events, goals and reactions. At this stage, each node and arc are of equal importance. The result from PUGG is the following plot unit graph:

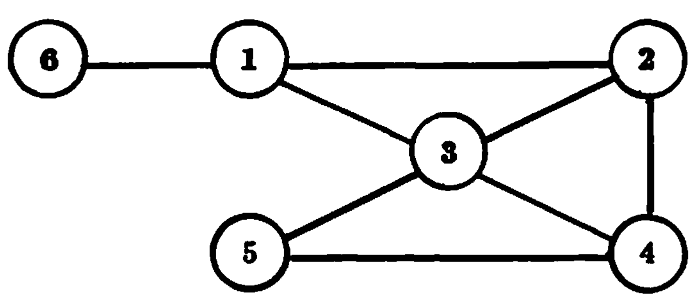

1. Motivation
2. Motivated Success
3. Change of Mind
4. Enabled Success
5. Nested Subgoals

Here we have a hierarchy of concepts (represented by nodes in the graph.) The most important node can be determined purely from the structure of the graph. Since this is a simple cluster we can use the summarization algorithm presented in section 3.3. The most important node for our summary is therefore change of mind. PUGG has taken an unranked, chronological list of states and events as input and produced a ranked set of plot units as output for our summary generation.

But how do we actually get that summary from this stage? Again, since PUGG deals only with the conceptual stage of this process, we will have to hook up another processor to the output to handle the details of generating the English summary.

5.3 Generation of English Language Summaries

To this end we are developing an interface between PUGG and MUMBLE, a natural language generation program (McDonald, 1983). Since MUMBLE can produce the English sentences but cannot determine how they should be ordered or what concepts from PUGG's output should be included, we are creating a text planner named PRECIS to stand between these two facilities (Cook, Lehnert & McDonald, 1984).

In its role of text planner, PRECIS decides what should be mentioned and what should be left for the reader to infer. Thus PRECIS interfaces between the purely conceptual and the purely linguistic and must have knowledge both of plot units and linguistic considerations. Precis associates different planning rules with different plot units or plot unit combinations. One such rule for the competition plot unit might be: mention both characters, the goal and the winner. We know that by naming the winner the reader will also make the natural inference that the other player necessarily lost. Let's consider a typical example. Suppose Betty beats out Nick for a higher ranking job which they both wanted:

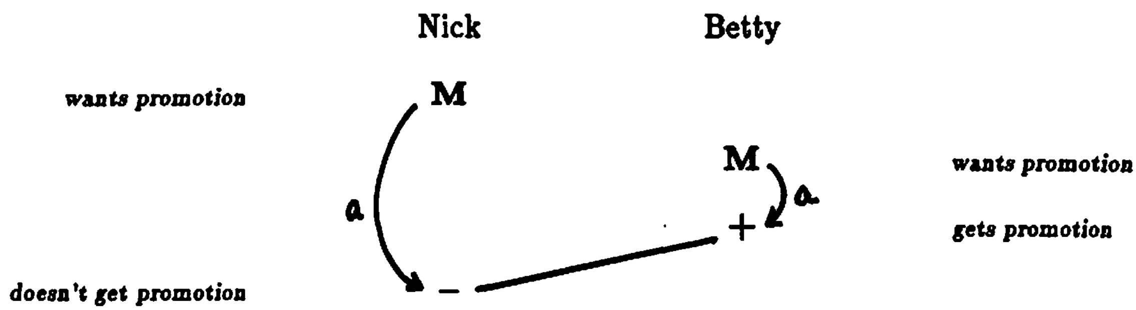

We could express this by stating "Nick wanted to get the promotion and Betty wanted to get the promotion, but Betty got it and Nick didn't." This sounds awkward and contains much more information than a reader needs. The sentence "Both Nick and Betty wanted the promotion, but Betty got it." is much more natural and is actually easier to understand. The reader is able to infer that since Betty got the promotion, Nick must not have, so we need not explicitly mention it. MUMBLE can handle such issues as pronominalization and the use of "both", but PRECIS must decide whether or not to explicitly include the loser.

PRECIS has been under development for only a short time and its stock of such rules is still small. We are currently developing more tactical rules and experimenting with the existing set to test the limits of their applicability. Eventually we will almost certainly need "meta-rules" to tell us which rule to choose when more than one might apply in a given situation.

5.4 Comparisons to Other Models

Several other systems for representing narrative structure or summarizing stories have been developed. Each takes a slightly different approach from the one we have taken here. A quick look at some of these models will help to point out their strengths and weaknesses, as well as some of the good and bad points of the plot unit representational system.

The only other extensive attempt to implement a computer model of text summarization was within the FRUMP system (DeJong, 1979). FRUMP analyzed UPI stories in about 50 domains, and provided summaries based on a top-down extraction of relevant information in those domains. The summaries were all based on an a priori set of expectations about the domain and did not exhibit much variation. For example, an earthquake story was summarized in terms of (1) where it occured, (2) what the Richter scale registered and (3) how many people were hurt. All earthquake summaries described these three components when available. This style of summarization was completely top-down and driven by specific expectations. FRUMP could not deal with unexpected information, and its summaries reflected total ignorance of anything unexpected, regardless of that information's importance.

With the plot unit approach, we do not form expectations until we have enough information to warrant their anticipation. In this way, unexpected information can be easily incorporated. FRUMP was specifically designed to handle news stories where, given a particular domain, the content is often quite predictable, and thus can work quite efficiently in such situations. Plot units are better suited to a less constrained input.

Another popular approach to the summary problem involves the notion of a story grammar (Rumelhart, 1977; Simmons & Correira, 1979; Thorndyke, 1977). Although many different story grammars have been proposed, the general idea is that stories are a linguistic form in much the same sense as sentences, and like sentences, can be described in terms of their constituent structure. In other words, just as we use a sentence grammar to identify the various parts of a sentence, we can use a story grammar to "parse" a story. Rummelhart points out that a number of short stories fall into what he calls the "EPISODE" schema. The EPISODE schema about protagonist **P** consists of:

1. Event **E** causes **P** to desire goal **G**

2. **P** tries to get **G** until outcome **O** occurs.

Each of the relational terms in this schema (cause, desire, and try) refer in turn to other schema that will likewise be instantiated by particular variables within a given story. Using these schema we can construct a hierarchical tree (with the EPISODE schema at its root) to represent the story.

Story grammar approaches to summarization try to anticipate these structural components, using them as a basis for forming a summary. Both story grammars and plot units can be used to predict human summarization behavior in terms of internal memory representations. Lehnert, Black and Reiser (1981) conducted a series of experiments comparing predictions of human summaries made by both the plot unit and story grammar approaches. The experimenters concluded that plot units predicted structural influences in the internal representation more effectively than story grammars. Story grammars have been criticized on many levels (Black & Wilensky, 1979) but the most basic limitation as a summarization model derives from being a purely top-down processor. As with FRUMP, story grammars are totally unable to handle input that does not conform to their expectations.

One theory of stories which seems to have much in common with the plot unit system is the theory of story points (Wilensky, 1982; 1983). Story points purport to capture the "storiness" of a narrative by representing those events in a story with high intrinsic interest. Such events are often "human dramatic situations" involving a problem for a character in the text. Although he does not propose the theory of story points as a model of summarization, Wilensky claims that a summary should consist of the point-related events of a narrative.

Unfortunately, Wilensky fails to provide us with a substantive strategy for text analysis with regard to story points. We are given no general method for determining what is interesting in a text. Wilensky characterizes the story point approach as being primarilly concerned with the affective reactions of the reader, while the plot unit approach deals with the emotional reactions of the story's characters, without considering the strong relationship between the two. It is reasonable to expect that a reader's response will be largely determined by the affective reactions (inferred or otherwise) of the characters in the story.

Both systems provide a high level memory representation for narratives. However, since the plot unit approach is motivated by the task of narrative summarization, it explicitly addresses the question of effective memory organization for efficient retrieval. This aspect of the problem has no counterpart in the points approach; Wilensky has nothing to say about the the problem of many points in competition for primary salience. His model must go one step further in this regard before it will be comparable to plot units.

5.5 Summary

We started by saying we wanted to learn more about understanding, both in machines and in people. Have we been successful? We have seen that understanding requires analysis of both the explicit and implicit content of a text. This analysis must occur at many levels ranging from descriptions of simple physical acts or events to the consideration of overall thematic patterns. These various levels seem to require different representational systems. Examining these representational schemes permits us to explore some of the possible inner workings of the understanding process.

Plot units provide one such system of representation at a fairly high level, since they are concerned with thematic patterns occuring in a text. We build these patterns up of smaller affective reactions and goal states. Plot units work well both as a computer program and as a cognitive model. PUGG is able to indentify the key concepts in a story and its output agrees well with summaries produced by human subjects.

We recall that the only real way we can evaluate understanding is to test for it through tasks like question answering or summarization. Applying the test of summarization to stories represented by the plot unit system gives us results very similar to human generated summaries. In this sense, then, we seem to have discovered one useful mechanism for both machine and human understanding. Although we cannot say that people use plot units exactly as described here in their internal processing of a text, the evidence is clear that we do employ some level of thematic processing in understanding narratives. For machines to fully understand, they too must employ such a level of processing.

Appendix A: Primitive Plot Units

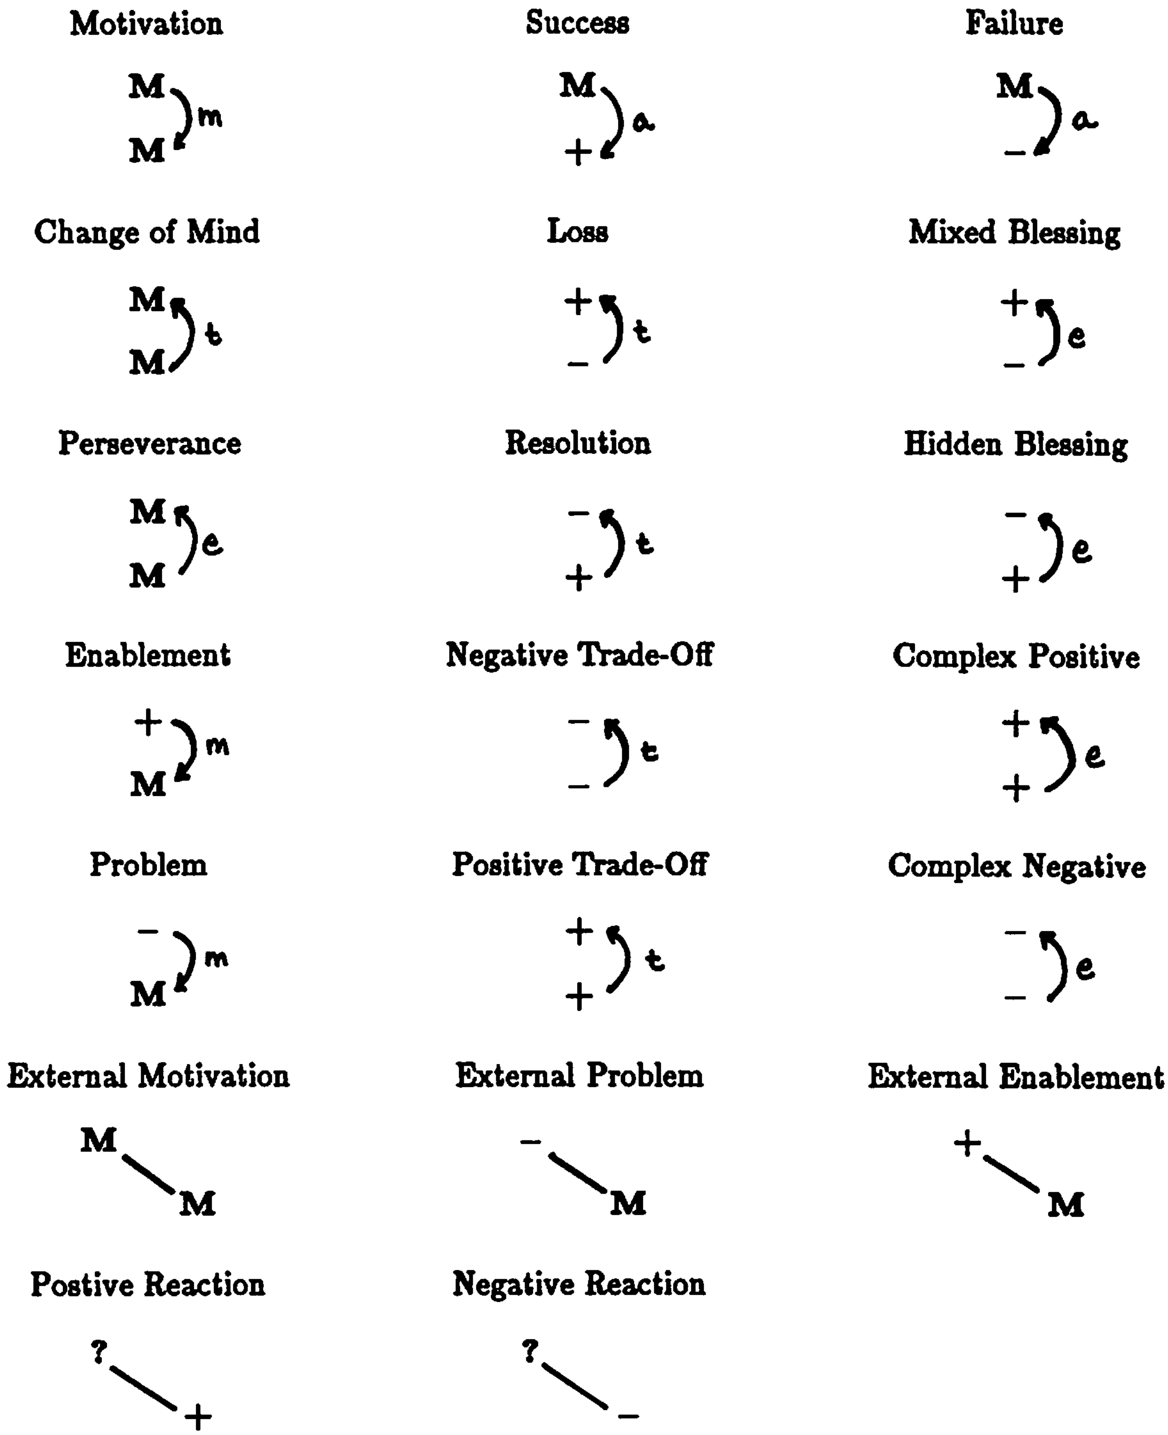

References

Black, J. B. and Wilensky, R. (1979) "An Evaluation of Story Grammars." *Cognitive Science* 3:3, pp. 213–230.

Brooks, JoAnn M. (1984) "Generation of Affect State Maps", *Proceedings of the First Annual Workshop on Theoretical Issues in Conceptual Information Processing*, Atlanta, GA.

Bruce, B., Collins, A., Rubin, A. D., and Gentner, D. (1980) "A Cognitive Science Approach to Writing" in Frederiksen, C. H., Whiteman, M. F., and Dominic, J. D. (Eds.) *Writing: The Nature, Development, and Teaching of Written Communication*, Hillsdale, N.J., Lawrence Erlbaum Associates, Inc.

Carbonell, J. G., Jr. (1978) "POLITICS: Automated Ideological Reasoning," *Cognitive Science*, 2:1 pp. 27–51

Cook, M., Lehnert, W. G., and McDonald, D. D. (1984) "Conveying Implicit Content in Narrative Summaries," paper to be presented at COLING-84, Palo Alto, California.

DeJong, G. F. (1979) "Skimming Stories in Real Time: An Experiment in Integrated Understanding" Department of Computer Science Research Report # 158, Yale University, New Haven, Connecticut.

Dyer, Michael G. (1981). "The Role of TAU's in Narratives", *Proceedings of the Third Annual Conference of the Cognitive Science Society*, Berkeley, Ca.

Dyer, Michael G. (1983). *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension.*, MIT Press, Cambridge, MA.

Fillmore, C. (1968) "The Case for Case," in Bach, E. and Harms, R. (Eds.), *Universals in Linguistic Theory*, New York, Holt Rhinehart & Winston.

Gee, James Paul and Grosjean, Francois (1984) "Empirical Evidence for Narrative Structure", *Cognitive Science*, 8:1, pp. 59–85.

Graesser, A. C. (1981) *Prose Comprehension Beyond the Word*, New York, Springer Verlag.

Hayes, J. R. and Flower, R. S. (1980) "Identifying the Organization of Writing Processes." in Gregg, L. and Sternberg, E. R. (Eds.) *Cognitive Processes in Writing*, Hillsdale, N.J., Lawrence Erlbaum Associates, Inc.

Kintsch, W. (1974) *The Representation of Meaning in Memory*, Hillsdale, N.J., Lawrence Erlbaum Associates, Inc.

Kolodner, J. L., (1983) "Reconstructive Memory: A Computer Model" *Cognitive Science*, 7:4, pp. 281–328.

Lehnert, Wendy G. (1979) "Representing Physical Objects in Memory" in Ringle, M. (Ed.) *Philosophical Perspectives in A. I.*, New Jersey, Humanities Press.

Lehnert, Wendy G. (1981). "Affect and Memory Representation", *Proceedings of the Third Annual Conference of the Cognitive Science Society*, Berkeley, Ca.

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

Lehnert, W. G., Alker, H., and Schneider, D. (1983) "Affective Plot Structure of Toynbee's *Christus Patiens*" in *Proceedings of the Sixth International Conference on Computers and the Humanities*. Washington, D.C.

Lehnert, W. G., Black, J. and Reiser, B. (1981) "Summarizing Narratives", *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Vancouver, British Columbia.

Lehnert, Wendy G., Dyer, Michael G., Johnson, Peter N., Yang, C. J. and Harley, Steve. (1983). "BORIS—An Experiment in In-Depth Understanding of Narratives", *Artificial Intelligence*, **20**, pp. 15–62.

Lehnert, W. G. and Loiselle, C. L. (1983) "Plot Unit Recognition for Narratives." Tecnical Report # 83-39, Department of Computer and Information Science, University of Massachusetts, Amherst, Ma.

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

Reiser, Brain J., Black, John B., and Lehnert, Wendy G. (1983) "Thematic Knowledge Structures in the Understanding and Generation of Narratives", *Discourse Processes*.

Reiser, Brian J., Lehnert, Wendy G., and Black, John B. (1981) "Recognizing Thematic Units in Narratives", *Proceedings of the Third Annual Conference of the Cognitive Science Society*, Berkeley, Ca.

Rieger, C. (1975) "Conceptual Memory and Inference," in Schank, R. *Conceptual Information Processing*, Amsterdam, North-Holland Publishing Co. Inc.

Roseman, Ira. (1979) "Cognitive Aspects of Emotion and Emotional Behavior", presented at the meeting of the American Psychological Association, New York. (unpublished manuscript, Dept. of Psychology, Yale University, New Haven, Ct.)

Rumelhart, D. E. (1977) "Understanding and Summarizing Brief Stories" in Laberge, D. and Samuels, S. (Eds.) *Basic Processing in Reading, Perception, and Comprehension*, Hillsdale, N.J., Lawrence Erlbaum Associates, Inc.

Rumelhart, D., and Norman, D. (1975) (Eds.) *Explorations in Cognition*, W. H. Freeman and Co., San Francisco, Ca.

Schank, R. C. (1975) *Conceptual Infromation Processing*, Amsterdam, North-Holland Publishing Co., Inc.

Schank, R. C. (1979). "Reminding and Memory Organization: An Introduction to MOPs", Dept. of Computer Science Research Report 170, Yale University, New Haven, Ct.

Schank, R. C. (1982). *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* Cambridge University Press, New York.

Schank, R. C., and Abelson, R. P. (1977) *Scripts, Plans, Goals, and Understanding.* Hillsdale, N.J., Lawrence Erlbaum Associates, Inc.

Simmons, R. F., and Correira, A. (1979) "Rule Forms for Verse, Sentences and Story Trees" in Findler, N. (Ed.) *Associative Networks—Representation and Use of Knowledge by Computers*, New York, N.Y., Academic Press.

Thorndyke, P. W. (1977) "Cognitive Structures in Comprehension and Memory of Narrative Discourse", *Cognitive Psychology* **9**, pp. 77–110.

Wilensky, R. (1980) *Understanding Goal-Based Stories*, New York, N.Y., Garland Publishers.

Wilensky, R. (1982) "Points: A Theory of the Structure of Stories in Memory" in *Strategies for Natural Language Processing*, Lehnert, W. and Ringle, M. H. (Eds.), Hillsdale, NJ, Lawrence Erlbaum Associates, Inc.

Wilensky, R. (1983) "Story Grammars versus Story Points," *The Behavioral and Brain Sciences.* **6**:4, pp. 579–623.

Wilks, Y. (1978) "Good and Bad Arguments for Semantic Primitives," *Communication and Cognition* **10**, pp. 181–221.

Winograd, T. (1972) *Understanding Natural Language*, New York, N.Y., Academic Press.

Woods, W. (1970) "Transition Network Grammars for Natural Language Analysis," *Communications of the ACM*, **13**:10.