Story Telling as Planning and Learning

Michael Lebowitz
Department of Computer Science
Columbia University
New York, NY 10027

November, 1985

CUCS-202-85

Revision of 'Story Telling and Generalization' -
as listed on the 1985 TR list.

To appear in *Poetics*.

Story Telling as Planning and Learning

Michael Lebowitz
Department of Computer Science -- Columbia University
New York, NY 10027, U.S.A.
5 November 1985

Abstract

The generation of extended plots for melodramatic fiction is an interesting task for Artificial Intelligence research, one that requires the application of generalization techniques to carry out fully. UNIVERSE is a story-telling program that uses plan-like units, "plot fragments", to generate plot outlines. By using a rich library of plot fragments and a well-developed set of characters, UNIVERSE can create a wide range of plot outlines. In this paper, we illustrate how UNIVERSE's plot fragment library is used to create plot outlines and how it might be automatically extended using explanation-based generalization methods. Our methods are based on analysis of a television melodrama, including comparisons of similar stories.

1 Introduction

In (Lebowitz, 1984) we described how extended story telling of the sort used to create fictional serials, novel series, television melodrama ("soap operas"), and the like involves a wide range of interesting Artificial Intelligence problems. We introduced a prototype program to generate extended stories, UNIVERSE, concentrating on its ability to create realistic characters. Here we describe a scheme for generating plot outlines that uses these characters along with plan-like units, "plot fragments". We show how UNIVERSE generates a simple piece of a plot outline and then discuss how the appropriate way to extend the capabilities of the program is for it to automatically expand its knowledge base using techniques closely allied to explanation-based learning methods, (DeJong, 1983; Mooney and DeJong, 1985), for example. The development of these learning methods has been based on analyses of plot outlines from a television melodrama.

The generation of extended stories -- those that continue on with no fixed end point -- raises a number of interesting issues in Artificial Intelligence. These issues include natural language generation,

knowledge representation, particularly about people and their character traits, knowledge-state assessment, memory organization and access, planning plot interactions, and so forth. We can expect research into extended story generation to both give us insight into the creative mechanism and lead to interesting story-telling programs, including systems where users weave themselves or characters they create and control into the story. Such programs would be valuable for both education and entertainment applications.

The story domain we have selected for UNIVERSE is that of interpersonal melodrama, as such stories provide excellent examples of narrative construction that are formulaic enough to generate by computer, and yet interesting enough to hold the interest of readers and viewers. Such a domain should help us study some of the simpler components of creativity. Eventually, we expect UNIVERSE to be able to generate connected stories in natural language form over a long period of time. For the moment, we are concentrating on generating plot outlines, and leaving problems of dialogue and other low-level text generation for later. As an illustration of the kind of plot outlines we would like to generate, consider the following synopsis of events from the U. S. television melodrama, *Days of Our Lives*:

> **STORY1** - Liz was married to Tony. Neither loved the other, and, indeed, Liz was in love with Neil. However, unknown to either Tony or Neil, Stephano, Tony's father, who wanted Liz to produce a grandson for him, threatened Liz that if she left Tony, he would kill Neil. Convinced that he was serious by a bomb that exploded near Neil, Liz told Neil that she did not love him, that she was still in love with Tony, and that he should forget about her. Eventually, Neil was convinced and he married Marie. Later, when Liz was finally free from Tony (because Stephano had died), Neil was not free to marry her and their troubles went on.

STORY1 exemplifies the kind of plot outline we would like UNIVERSE to generate at this point in our research as a precursor to full story generation. There are several important points to note about this example. First, the interactions among characters are quite complex. It is important that the behavior of all the characters make sense in terms of what we already know about them. Actions must not violate rules of the world (be *consistent*) and must be motivated (be *coherent*). On the other hand, it is not enough that we simply simulate the lives of these characters, as was done in TALE-SPIN (Meehan,

1976), since it is unlikely that people would naturally act in ways as interesting as in STORY1 or that their actions would interleave with each other so nicely. Our generation scheme and proposal for generalizing new plot fragments both endeavor to create plot outlines that are believable and yet interesting.

2 The basic UNIVERSE story-telling model

Story telling in UNIVERSE is a plan-based activity. It is based around a library of "plot fragments" that play the same role as standard plans in planning systems such as those of Sacerdoti (1977), Schank and Abelson (1977), and Wilensky (1983). We will discuss in Section 3 how the library of plot fragments might be created or extended automatically. The plot fragments provide narrative methods that achieve goals. What is crucial here is that the goals and plot fragments used by UNIVERSE should not be viewed as goals and plans of the characters (although these must be monitored to maintain consistency), but, instead, goals and plans of the *author* (or program, in our case). This allows the kinds of actions by characters that make sense, and yet are not necessarily what a character would choose to do if an independent agent. See (Dehn, 1981) for additional reasons for using author goals. This approach fits nicely with work in narrative theory such as that of Barthes (1977), Todorov (1977) and Eco (1979), which has influenced the development of UNIVERSE, as has the Artificial Intelligence work of Meehan (1976), Dehn (1981), and Yazdani (1983); Turner and Dyer (1985) is also related.

As a system that uses plans to generate language, UNIVERSE is related to work that has been done in low-level text generation, including that of McDonald (1980), McGuire (1980), Appelt (1982), McDonald and Conklin (1982), and McKeown (1982). However, UNIVERSE operates at a much more conceptual level (although eventually we will, of course, also need to generate text). It is important to remember that an entire UNIVERSE story cannot be planned out, since the kinds of stories we envision theoretically do not end. Instead, we tell the story as we plan (which means the opportunity for backup is limited). Story telling becomes the expansion of goals, in the problem-solving sense, until they lead to actual events that can be generated.

The plot fragments used by UNIVERSE can span a wide range of levels, from very general, thematic plans that may unfold over a long time in the story, to plans for specific actions. They may also include goals whose sole purpose is setting the stage for later events (e.g., "dropping hints"). The more general plot fragments are much like the plot units of Lehnert (1981) which describe stories in terms of abstract emotional impact. For example, a plot fragment might be to cause an occurrence of John "double crossing" Mary (perhaps to achieve the author goal of having Mary dislike John). The expansion of this plot fragment would require the creation of many other goals, the precise nature of which would depend on what we know about John and Mary. For example, if Mary was a stock broker, John might double cross her by giving her phony "inside stock market information".

Before looking in detail at a typical plot fragment, it is worth mentioning one important aspect of our methods. UNIVERSE assumes that it has available an interesting and coherent set of characters to work with. When starting a new story, the system builds up a story-telling universe of characters as described in (Lebowitz, 1984). This is precisely the way that some authors work in developing novels: create a set of characters and work from there. For example, Eco (1984) says that "What I mean is that to tell a story you must first of all construct a world, down to the slightest details" (page 2-3). He goes on to describe how this "world" must include "lists of names and personal data for many characters" (page 2-4). Figure 1 shows a typical UNIVERSE character.

UNIVERSE representations of characters contain quite a bit of information. For purposes of this paper, the important parts for each character are a set of traits that describe the person and a set of interpersonal relationships (IPRs) involving the character. The IPRs are broken into numeric scales that are not relevant here. In addition, there are stereotypes that "explain" (motivate) the traits and interpersonal relationships to provide coherence. In Figure 1, Liz Chandler is basically a rich, intelligent socialite who is currently married to Tony Dimera, but hates him (the last point begin recorded in the

```
Name: LIZ CHANDLER (LIZ)
Marriages:
   DON CRAIG [DON] [&MF1] [1980]
   TONY DIMERA [TONY] [&MF3]
IPRs:
   HUSBAND-WIFE     TONY DIMERA [TONY]            0/-8//8/8//6/-6//7/-3
   EX-SPOUSES       DON CRAIG [DON]               -5/-5//4/4//0/0//4/4
Stereotypes: ACTOR KNOCKOUT SOCIALITE PARTY-GOER
Trait modifiers: (SEX F) (AGE YA) (WEALTH 3) (PROMISCUITY -3)
                 (INTELLIGENCE 3)
Overall description:
WEALTH              8
PROMISCUITY         3
COMPETENCE          NIL
NICENESS            0
SELF-CONF           6
GUILE               7
NAIVETE             7
MOODINESS           6
PHYS-ATT            7
INTELLIGENCE        7
GOALS               (FIND-HAPPINESS BECOME-FAMOUS MEET-FAMOUS-PEOPLE)
AGE                 YA
SEX                 F
```

Figure 1: A typical UNIVERSE character representation

numbers following their interpersonal relationship).[1] UNIVERSE also keeps track of some goals the character might have, for use in plot generation. Further details of character representation can be found in (Lebowitz, 1984).

Having available a rich set of characters greatly improves UNIVERSE's ability to combine plot fragments as well as diversifies the kinds of stories it can tell. This is superior to creating characters on the fly for reasons of coherence. We now return to looking at the plot fragments that use these characters. Figure 2 shows the main features of a typical UNIVERSE plot fragment, *forced-marriage.* It involves an evil parent forcing one his children to remain in an unhappy marriage, preventing that child from being with the person they really love. It will form the core of UNIVERSE's generation of a simplified version of STORY1.

---

[1]These characters are adapted from *Days of Our Lives*, although UNIVERSE is able to generate equally interesting characters itself as described in (Lebowitz, 1984).

```
PLOT FRAGMENT: forced-marriage

CHARACTERS:  ?him ?her ?husband ?parent

CONSTRAINTS: (has-husband ?her)          {the husband character}
             (has-parent ?husband)       {the parent character}
             (< (trait-value ?parent 'niceness) -5)
             (female-adult ?her)
             (male-adult ?him)

GOALS: (churn ?him ?her)  {prevent them from being happy}

SUBGOALS: (do-threaten ?parent ?her "forget it") {threaten ?her}
          (dump-lover ?her ?him) {have ?her dump ?him}
          (worry-about ?him)     {have someone worry about ?him}
          (together * ?him)      {get ?him involved with someone else}
          (eliminate ?parent)    {get rid of ?parent (breaking threat)}
          (do-divorce ?him ?her) {end the unhappy marriage}
          (or (churn ?him ?her)     {either keep churning or}
              (together ?her ?him)) {try and get ?her and ?him back together}
```

Figure 2: A typical UNIVERSE plot fragment

The first piece of information in Figure 2 is a list of the roles involved in the plot fragment. Roles are designated with "?"s. Here we have the lovers (?him and ?her) along with the woman's husband and the husband's parent. Along with each role can be a stereotype that the character must fit. Here these do not appear, as any class of character may be used for each role, but often we may have a role that can only be filled by a certain type of character, such as a doctor in a malpractice plot fragment or a lawyer in a divorce case.

Following the roles is a list of additional constraints that the characters must satisfy. In this case, most of the constraints simply verify the family relations required. However, we also require that the "evil parent" be a nasty sort -- have a low "niceness value"; see (Lebowitz, 1984) for more about traits. This is necessary for the plot fragment to be at all believable. For the roles that are not specified by the goal, UNIVERSE will attempt to find characters that fit the constraints (or create them, if need be).

The next piece of information about the *forced-marriage* plot fragment is the goals that it can be used to achieve. Here the goal is churn -- keep two lovers apart. Obviously this goal makes no sense from the point of view of the characters involved, but it makes a great deal of sense for the author, and,

indeed, is a staple of melodrama ("happily ever after" being notoriously boring in fiction, if not in life). UNIVERSE has a number of other plot fragments for achieving this goal, such as lover's fights and job problems. Although *forced-marriage* is used for only one goal, it is possible for plot fragments to serve several goals, a fact which plays an important role in our generation algorithm.

Finally, Figure 2 shows the heart of the plot fragment -- an ordered series of subgoals that must be achieved to carry out the plot fragment (actually, any partial ordering of the subgoals may be specified). Often a plot fragment will include actually text to be generated, but since *forced-marriage* is a relatively high level plot fragment, it simply spins off subgoals. It causes the parent to threaten his daughter-in-law, which leads her to break off her love affair. After being consoled by some sympathetic character, her lover then gets involved with someone else. Again, this would not necessarily be the character's goal at the time, but the author wants this to make things more difficult later on. Eventually, the parent is eliminated, in one way or another, which removes the threat. The woman can then divorce her husband. Finally, the plot fragment either churns the relationship further, or tries to get the estranged couple together (which may be pleasingly complicated, as he is now presumably involved with someone else). Each of the subgoals in Figure 2 can potentially be satisfied by a variety of different plot fragments, leading to a wide range of possible stories.

Additionally, we can include in plot fragments information about a given fragment effects the characters involved. This means primarily the modification of character traits and interpersonal relations, causing a character to become cynical or one to distrust another, for instance. For *forced-marriage*, this mostly occurs in the subgoals. For example, the threat is likely to make the women strongly dislike her father-in-law (if she did not already).

With a library of plot fragments and a set of characters UNIVERSE is ready to tell stories. The basic story-telling algorithm is a straightforward one, relying mainly on the richness of its plot fragment and character databases. UNIVERSE maintains a precedence graph that records how the various

pending author goals and plot fragments relate to each other and to events that have been told already. Then, the program selects an author goal to expand, and continues this process recursively, until enough goals reach "ground level", i.e., actual events, at which point the events would be told to the reader using natural language generation methods. This algorithm, much like the one used for TALE-SPIN (except, to reiterate, that it uses author goals) is summarized in Figure 3.[2]

Pick a goal with no missing pre-conditions

Pick a plot fragment for that goal,
achieving extra goals, if possible

"Execute" the plan, including
adding new goals to the goal graph and
"telling" (producing output), if appropriate

Figure 3: The basic UNIVERSE story-telling algorithm

One of the nice side effects of the algorithm in Figure 3 is that as multiple goals are pursued, various "plot lines", i.e., the pursuit of different high level plot fragments, will become interleaved. This, along with the plot library and universe of characters leads to the production of quite intricate plot outlines. Before we discuss a few additional details of the story-telling algorithm, we will look at the production of a simple plot outline based on the *forced-marriage* plot fragment. The UNIVERSE output is shown in Figure 4.

Basically, the program trace in Figure 4 shows the expansion of the *forced-marriage* plot fragment for the churn goal. The input is a request to achieve various goals, a request that could have been generated by the program itself in achieving higher level goals. Once UNIVERSE has selected this plot fragment to use in pursuit of the input goal, each of the subgoals is pursued in turn. In the UNIVERSE's library at the time this plot outline was generated, each subgoal could be achieved by a single plot fragment. In a full-blown system, of course, each of these subgoals could lead to a range of subgoals,

[2]Actually, the algorithm is even more similar to the micro-TALE-SPIN program described in (Charniak et al., 1980).

```
*(tell '(((churn neil liz))))

working on goal -- (CHURN NEIL LIZ)
Several plans to choose from FORCED-MARRIAGE LOVERS-FIGHT JOB-PROBLEM
-- using plan FORCED-MARRIAGE

working on goal -- (DO-THREATEN STEPHANO LIZ "forget it")
-- using plan THREATEN

>>> STEPHANO threatens LIZ: "forget it"

working on goal -- (DUMP-LOVER LIZ NEIL) -- using plan BREAK-UP

>>> LIZ tells NEIL she doesn't love him

working on goal -- (WORRY-ABOUT NEIL) -- using plan BE-CONCERNED
Possible candidates -- MARLENA JULIE DOUG ROMAN DON CHRIS KAYLA
Using MARLENA for WORRIER

>>> MARLENA is worried about NEIL

working on goal -- (TOGETHER * NEIL)
Several plans to choose from SEDUCTION DRUNKEN-SNEAK-IN
SYMPATHETIC-UNION JOB-TOGETHER
-- using plan SEDUCTION
Possible candidates -- DAPHNE RENEE
Using DAPHNE for SEDUCER

>>> DAPHNE seduces NEIL

working on goal -- (ELIMINATE STEPHANO)
Several plans to choose from ATTEMPTED-MURDER EXPOSE
-- using plan ATTEMPTED-MURDER
Possible candidates -- RENEE ALEX
Using RENEE for KILLER

>>> RENEE tries to kill STEPHANO

working on goal -- (DO-DIVORCE TONY LIZ) -- using plan DIVORCE

>>> LIZ and TONY got divorced

working on goal -- (TOGETHER LIZ NEIL)
no acceptable plans
```

Figure 4: Creating a simple plot outline

which in turn might have to be expanded further. In Figure 4, ">>>" indicates the final output of the program, that is, the actual plot outline. The remainder of the output describe the various goals and plans that are being considered and pursued.

UNIVERSE can handle a number of interesting cases that go beyond this example. The first

comes in picking the goal to be pursued. The program does not simply pursue goals in a depth-first manner. We allow the specification of goals that may be pursued in parallel, which are added to the precedence graph. Since fiction is an inherently serial form, on each cycle of the program, UNIVERSE must pick a single goal to pursue. However, since it does that by checking for goals with all prerequisites satisfied (rather than using a depth-first approach), UNIVERSE can perform interesting sorts of switches among the pursuit of different goals. We are experimenting with a variety of heuristics to handle the case where several goals could be validly pursued.

The desire to plan a story in an interesting and effective way also comes into play in the selection of a plan to use in the pursuit of a given goal. UNIVERSE makes use of a very opportunistic plot fragment selection method. When there are multiple plot fragments that can achieve the same goal, or multiple characters that can be used in a plot fragment, both of which occur with increasing frequency as we build up the program's plot fragment library, UNIVERSE tries to select the one that will achieve additional goals that the program has in its precedence graph. This is illustrated in Figure 5, where we give the program an additional goal -- to get Neil and Renee together -- as well as churning Neil and Liz's relationship.

In both Figure 4 and Figure 5, UNIVERSE picked the *seduction* plot fragment to achieve the goal of getting Neil involved with someone after Liz broke up with him. In Figure 4, the program simply picked a seductress at random from among the characters with the characteristics that fit the constraints of the plot fragment (had none existed, it would have created one). In Figure 5, however, UNIVERSE was able to satisfy its second goal as a side-effect of the seduction by selecting Renee as a seductress. A further result of this is that Renee cannot take part in the *attempted-murder* plot fragment, as she was involved in the earlier seduction. This sort of opportunistic planning goes along way towards achieving the intricate plot interconnections that exist in most popular melodrama.

```
*(tell '(((churn liz neil) (together renee neil))))

working on goal -- (CHURN LIZ NEIL)
Several plans to choose from FORCED-MARRIAGE LOVERS-FIGHT JOB-PROBLEM
-- using plan FORCED-MARRIAGE

working on goal -- (DO-THREATEN STEPHANO LIZ "forget it")
-- using plan THREATEN

>>> STEPHANO threatens LIZ: "forget it"

working on goal -- (DUMP-LOVER LIZ NEIL) -- using plan BREAK-UP

>>> LIZ tells NEIL she doesn't love him

working on goal -- (WORRY-ABOUT NEIL) -- using plan BE-CONCERNED
Possible candidates -- MARLENA JULIE DOUG ROMAN DON CHRIS KAYLA
Using MARLENA for WORRIER

>>> MARLENA is worried about NEIL

working on goal -- (TOGETHER * NEIL)
Several plans to choose from SEDUCTION DRUNKEN-SNEAK-IN
SYMPATHETIC-UNION JOB-TOGETHER
Possible candidates -- DAPHNE RENEE
Using RENEE for SEDUCER

>>> RENEE seduces NEIL

working on goal -- (ELIMINATE STEPHANO)
Several plans to choose from ATTEMPTED-MURDER EXPOSE
-- using plan ATTEMPTED-MURDER
Using ALEX for KILLER

>>> ALEX tries to kill STEPHANO

working on goal -- (DO-DIVORCE TONY LIZ) -- using plan DIVORCE

>>> LIZ and TONY got divorced

working on goal -- (TOGETHER LIZ NEIL)
no acceptable plans
```

Figure 5: Same situation with extra goal

3 The next step: Generalizing plot fragments

We feel the generation framework outlined in Section 2, which is able to generate moderately complicated and interesting plot outlines, provides a good basis for future work. To generate interesting stories over the long run, however, it will be necessary for the program to be able to automatically expand its plot fragment library by creating new plot fragments. We are studying this problem by looking at

various plot outlines from *Days of Our Lives*. For example, consider STORY2:

> **STORY2** - Hope and Bo were very much in love. However, Bo was involved in some dangerous activities and was worried about Hope's safety. To protect her, he told Hope that he didn't love her, that he was living with Diane, a childhood friend, and that she should get on with her life. She was more or less convinced, and started spending a lot of time with Larry, the District Attorney, who wanted to marry Hope for political purposes. Just as they were about to be married, Bo arrived, the danger having passed, and told Hope that he loved her. They ran off together. However, Maxwell, a new evil character, who was interested in Larry's career, had his goons capture Hope (without Larry's knowledge) and tell her that she was to return to Larry or harm would come to Bo and her other friends. She did so (after some of Maxwell's goons beat up Bo), and married Larry, convincing Bo she no longer cared for him. Then, Megan, Maxwell's daughter and a childhood love of Bo's appeared and told him she once bore his child. We assume that when Hope is finally free of Larry, Bo will be entangled with Megan.

STORY2 at first appears rather complex. However, if we look at it closely, we realize that it is really two sequential plot lines, each quite similar to STORY1, but with a few twists. This can be seen more clearly if we break down STORY1 and the two plot lines in STORY2 into sequences of events displayed in parallel. This is done in Figure 6.

Figure 6 makes clear the similarities among the three plot lines. For example, the first parts of STORY1 and STORY2b differ primarily in character substitutions. STORY1 and STORY2a, on the other hand, are similar at a somewhat more abstract level. So, for instance, the B events in STORY1 and STORY2b involve a marriage and an engagement, respectively, while STORY2a involves some sort of dangerous activity. To recognize the similarity among these three plot lines, we have to realize that each of these events plays the same role in their stories -- they create a reason to keep the lovers apart. The I events, in the same fashion, have rather different actions playing the same role -- to further thwart the lovers even when the prime obstacle between them disappears.

Notice that while these three plot lines are all quite similar, we cannot hope to generate STORY2a or STORY2b from the same rather specific plot fragment that we used to generate STORY1 (the *forced-marriage* plot fragment shown in Figure 2, or even a somewhat more complex version of it). *Forced-marriage* has rather specific constraints and calls for rather specific events -- too specific to

	STORY1	STORY2a	STORY2b
A	Liz loved Neil	Bo loved Hope	Hope loved Bo
B	Liz was married to Tony	Bo was involved in activity that could endanger Hope	Hope was engaged to marry Larry
C	Stephano wanted Liz married to Tony		Maxwell wanted Hope married to Larry
D	Stephano told Liz if she left Tony he would hurt Neil		Maxwell told Hope, indirectly, if she didn't marry Larry, he would hurt Bo
E	Stephano had Neil hurt in a bomb explosion		Maxwell had Bo beat up
F	Liz told Neil she didn't love him	Bo told Hope he didn't love her	Hope told Bo she didn't love him
G	Liz stayed with Tony	Bo lived with Diane	Hope married Larry
H	Neil believed her	Hope believed him	Bo believed her
I	Neil married Marie	Hope got involved with Larry and was about to marry him	Bo discovered Megan had had his child
J	Liz eventually got free of Stephano, but couldn't be with Neil because he was married to Marie		???
K	Still later, Liz was able to marry Neil	Bo's danger passed, and Bo and Hope ran off together	

Figure 6: Three similar plot outlines

accommodate STORY2a or STORY2b. For example, in STORY2b, instead of an overt threat from the woman's father-in-law there is an implicit threat from the danger Bo is in. We could, however, hope to derive from *forced-marriage* a more general plot fragment that would produce STORY2a or STORY2b.

Based on this sort of analysis of existing melodramatic plot outlines, we propose to automatically

augment UNIVERSE's plot fragment library by generalizing existing plot fragments.[3] In general terms, we can easily see how this could lead to the production of STORY2a and STORY2b. The *forced-marriage* plot fragment could be generalized into a "coerce into staying out of a relationship" plot fragment where step **B** involved a competing goal, step **D** a threat (optional), step **G** a competing relationship and so forth. Then this plot fragment could be instantiated into the specific stories STORY2a or STORY2b (or back into STORY1).

The main process of generalization proposed here is simply a relaxation of the constraints upon the role fillers for a given plot fragment along with modification of the remainder of the plot fragment to accept any role filler that fits the generalized constraints. So, for example, instead of requiring a married couple in *forced-marriage*, we might have UNIVERSE just look for any sort of a positive emotional relationship. Various details will be added to the actions and subgoals depending on how the plot fragment is instantiated. A similar process might involve generalizing the subgoals required in the plot fragment, replacing them with more abstract goals.

While it is easy to simply generalize the role constraints (as described below, we have implemented such generalization), it is obviously not acceptable for such generalization to be arbitrary (or else we would just generalize all role fillers to "person"). We must consider two issues: 1) keeping the generalized plot fragments believable (consistent and coherent) and 2) making sure the plot fragments still generate interesting stories.

We feel that the maintenance of believability will be accomplished through processing much like the explanation-based learning methods of DeJong (1983) and Mooney and DeJong (1985).[4] The basic

---

[3]An alternate, and ultimately quite similar, method would be to tell new stories by direct modification of stories told previously, without going through generalizing plot fragments -- essentially by analogy. Generalizing new plot fragments fits better into the UNIVERSE story-generation scheme.

[4]There has been considerable other work into explanation-based learning, e.g., (Carbonell, 1983; Mitchell, 1983; Mostow, 1983; Ellman, 1985); also see (Michalski et al., 1986). The story understanding domain of DeJong and Mooney is most closely related to our work.

idea is that we would build up a causal analysis of a plot fragment, using methods such as those described in (Schank and Abelson, 1977; Schank, 1982), and then generalize it in various ways such that the analysis still holds, as done in (DeJong, 1983; Mooney and DeJong, 1985), to build up new story understanding schemata. Using the causal analysis, we can determine various ways that a plot fragment can be generalized while maintaining that same causal explanation. So, for example, in the *forced-marriage* plot fragment, we might be able to generalize the relation that one participant is being forced into from a marriage to any permanent relationship (e.g., "living together"), but not into a less personal relationship, like ("roommates"), as that would undermine the causal explanation for the plot fragment.

There are two primary ways in which our generalization methods differ from the explanation-based learning of DeJong and Mooney. Both of these differences address the second of our concerns, that the generalized plot fragments still yield interesting stories. We propose that: 1) author goals be included in the explanations of the plot fragments and 2) plot fragments not be generalized as far as possible all at once.

Inclusion of author goals in the causal analysis is crucial. In many cases, it will be impossible to analyze a plot fragment without considering such goals. For example, in *forced-marriage*, the only rationale for getting the jilted lover involved with another person is to satisfy an author goal of complicating the situation ("churning" it). The characters would certainly not want this to happen, particularly if they knew the original relationship would become feasible again. Even when it is possible to analyze a plot fragment based on character goals, it is likely that only the consideration of author goals will inhibit over-generalization. So, continuing our *forced-marriage* example, we might imagine a generalization that makes real-world sense if the main characters merely like, not love, each other. However, such a plot fragment is unlikely to produce interesting stories, as it will not further author goals that revolve around intense relationships.

Our other point is that we do not always want to generalize plot fragments as far as possible, even within the constraints of the explanation. It is unlikely that we will be able to analyze plot fragments in enough detail to fully understand the "flavor" that they provide. We will maintain the flavor of the initial fragments by only generalizing some of the features of each plot fragment at a time. This will help both by simplifying the instantiation process (as there will be fewer degrees of freedom for the various role fillers) and will maintain most of the interesting aspects of the plot fragments. The richness of the stories that we are looking for comes from instantiating the generalized plot fragments with a variety of different role fillers. The choice of each role filler can then lead to interesting changes required to keep the story consistent (the determination of which can, incidentally, also make use of the causal analysis needed for generalization).

Even small generalizations of plot fragments will add considerable richness to UNIVERSE's stories. Discovering that *forced-marriage* "works" if the outside threat comes to achieve political goals rather than desire for a grandchild, or that it will work if the threat is made to the man, not the woman, will expand the utility of the plot fragment. Such generalization is well within the scope of current explanation-based learning methods.

In order to further illustrate the generalization methods we are suggesting, we will consider how the plot fragment we have been looking at, *forced-marriage*, could be generalized. We will sketch the sort of explanation that might be built up for this plot fragment (though not in the full level of detail that would ultimately be necessary) and show how the explanation would affect the allowable generalizations.

Before laying out an explanation for *forced-marriage*, an explanation that will relate the events in the fragment to the goals it achieves, we must look briefly at the relevant author goals. The most most important such goal to consider involves "churn", the top-level goal used in *forced-marriage*. Churn is naturally decomposed into two components for analysis purposes. The first is MAINTAIN-ROMANTIC-TENSION (MRT) -- establishing a state wherein two characters are in love and one is free to pursue the

relationship between them, but the other is not. We define this goal as the set of states show in Figure 7. This goal state is simply assumed to be an interesting one with no further justification other than our own experience with melodramatic stories. This avoids the need for detailed analysis of what makes a plot fragment interesting.

```
MAINTAIN-ROMANTIC-TENSION (MRT):
  in-love (A, B)
  available (A)
  not available (B)
where
  not available (A) if and only if married (A, X) or engaged (A, X) for some X
```

Figure 7: A possible author goal

The definition of MRT introduces one problem with respect to the churn goal. Since the initial state needed for churn is the one desired by MRT, there must be another active goal to justify perturbing this state of affairs. This is dealt with by the second aspect of churn, a desire to KEEP-THE-STORY-MOVING (KTSM) which basically captures the idea that having the characters do *something* is essential, and that even if the effect is not optimal, we can always restore an interesting state of the world.

With the two relevant (and somewhat conflicting) author goals, MRT and KTSM, we can look at an analysis of *forced-marriage*. A simplified form of such analysis is shown in Figure 8. The notation for character goals in Figure 8 is that each goal is assumed to be held by the character in brackets. So, together(?him, ?her) [?him] means that ?him has the goal of getting ?him and ?her together. In this case, ?her has the same goal, but that need not be the case. The arrows in Figure 8 indicate some unspecified form of causal connection.

Figure 8 tries to capture the reasons why the various subgoals in the plot fragment are included in the plot fragment, their motivation, if you like. The analysis basically breaks into two parts -- the initial

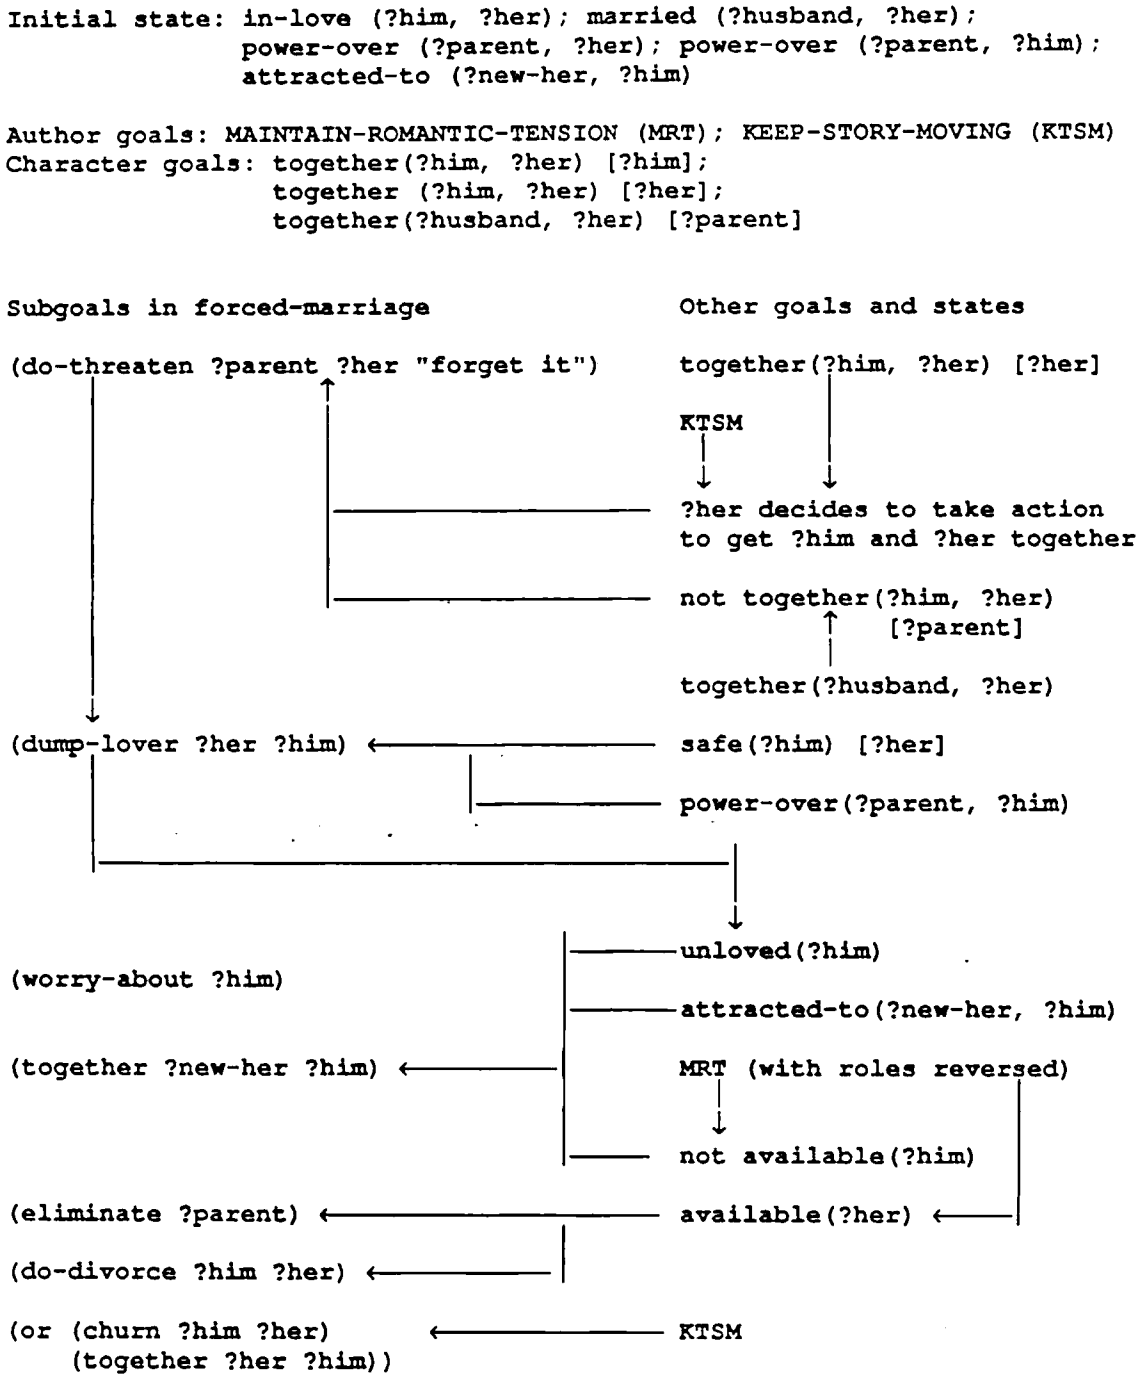

Figure 8: Rough analysis of *forced-marriage*

threat along with its effects and the re-establishment of romantic tension. We see that the threat's occurrence makes sense only if the woman had decided to take action to pursue the love affair. This, along with the parent's power and his desire to maintain the existing marriage motivates the threat.

(Elements of various character traits might also come into play.) The woman's desire to keep her lover safe, along with the parent's power motivate her forcing a breakup.

Once the relationship has been broken (at least from one character's perspective), there is a need to pursue a new author goal (since the romantic tension is broken). One possibility is simply to re-establish MRT, but with the roles reversed (i.e., the other partner being available), and that is what happens here. The new MRT goal motivates both the steps needed to get the man involved with someone else and to end the woman's marriage (which must include the removal of the threat over her and her lover). Note that to be believable, these goals must be pursued in this order, or the lovers would simply get together after the woman's marriage dissolved. In a fully developed *forced-marriage*, we would probably need another step to make sure that the in-love goal again held, as it was presumably broken by the breakup.

Our point here is not the details of this analysis, but rather how the analysis effect the ways we can generalize *forced-marriage.* Specifically, it restricts how far we can generalize the subgoals that are pursued and the constraints on the characters. We shall look at a few examples of this.

Our analysis of *forced-marriage* does not severely restrict the generalization of constraints on the characters. The lovers could basically be any pair of people with a romantic attraction. The husband role need not literally be the woman's husband -- any relation that is threatened by the desired relationship of the lovers will do. The one place we do get strong constraints is with the parent. While this character does not have to actually be the husband's parent, it does have to be someone with both an interest in maintaining the relation between the husband and the woman and also have enough power to successfully threaten the woman. A more detailed analysis might add additional restrictions about personality traits of characters in *forced-marriage*, e.g., that the woman not be so headstrong that she would ignore the parent's threat, and the parent must be sufficiently evil for the threat to be taken seriously.

The causal analysis in Figure 8 also restricts how we might generalize the subgoals pursued in *forced-marriage*. For example, since the goal of getting the man involved with someone new is just a means to achieve the higher-level goal of making him not available to the woman, we can generalize the together goal into a set of goals that achieve the desired purpose. For example, the plot fragment would make sense if we had the man get involved in a dangerous activity or contract amnesia -- anything that would inhibit his getting back with the woman. Similarly, the goal of eliminating the parent can be replaced by a more general one of removing his ability to threaten the lovers.

Generalizing *forced-marriage* while adhering to these restrictions would allow us to develop a plot fragment that covers a number of "forced relationship" situations. Despite being a single plot fragment, it could be instantiated in a variety of ways depending on the characters involved -- business blackmail, a homosexual relationship, and so -- and yet still maintain the aspects of *forced-marriage* that cause it to make sense and be interesting.

There is one further interesting issue involving Figure 8. We note that the subgoal of having someone worry about the man does not play a part in our causal analysis. In using explanation-based learning for most tasks, the absence of connections of an element would cause it to be deleted from a generalized description, on the theory that it must not be important. Here, however, we take the opposite approach. We make the assumption that although we do not understand why this subgoal is important, it may well play a significant role in making the fragment "work" (although not in causing it to make sense), perhaps carrying out an unknown author goal. Thus, we should leave it alone and include it in the generalized plot fragment. Only detailed computer experiments will determine whether this is the correct approach.

Finally, it is worth reiterating here that the analysis in Figure 8 is considerably simplified from what would ultimately be needed. Among the ways the analysis would have to be made more complex are: 1) analyzing at a lower level of detail, such as Conceptual Dependency (Schank, 1972); 2) further classifying

the various causal links, as in (Schank, 1975); 3) simply including more character and author goals. The first two points were basically done in (DeJong, 1983; Wilensky, 1983; Mooney and DeJong, 1985).

To illustrate the utility of generalizing plot fragments, we have developed a simple UNIVERSE module to generalize constraints but without the analysis of the underlying causality just described. Figure 9 shows the generalized constraints for *forced-marriage*. Basically, we have built the causality restrictions into the allowed role generalizations. We can see in Figure 9 that the requirement for the female lover to have a husband with a nasty parent has been generalized into a requirement for one lover to have a spouse with a nasty parent. The requirement for opposite sex lovers has also been relaxed. This gives us a plot fragment that can describe pressure on either side of a opposite or same sex relationship (given proper definition of the predicate "has-spouse"). Note that the names of the role fillers (e.g., ?him, ?her) are now rather misleading. The level of nastiness needed by the parent has also been relaxed.

```
CONSTRAINTS: (has-spouse ?her)          {instead of has-husband}
             (has-parent ?husband)
             (< (trait-val ?parent 'niceness) 0)
             (adult ?her)               {instead of female-adult}
             (adult ?him)               {instead of male-adult}
```

Figure 9: Generalized constraints for *forced-marriage*

Figure 10 shows how the generalized version of *forced-marriage*, *forced-marriage0*, can be used by UNIVERSE. When we ask the program to "churn" a relationship between David and Renee, it can use the generalized plot fragment, since David's wife's father (Alex) is nasty enough to threaten David. The original *forced-marriage* would not be applicable since Renee does not have a nasty father-in-law. After the selection of *forced-marriage0*, processing proceeds much as in Figure 4. It is worth reiterating that the plausibility of *forced-marriage0* is somewhat fortuitous, since we have not implemented the causal analysis that will ultimately be necessary.

```
*(tell '(((churn david renee))))

working on goal -- (CHURN DAVID RENEE)
-- using plan FORCED-MARRIAGE0 (FORCED-MARRIAGE with generalized constraints)

working on goal -- (DO-THREATEN ALEX DAVID "forget it") -- using plan THREATEN

>>> ALEX threatens DAVID: "forget it"

working on goal -- (DUMP-LOVER DAVID RENEE) -- using plan BREAK-UP

>>> DAVID tells RENEE he doesn't love her

working on goal -- (WORRY-ABOUT RENEE) -- using plan BE-CONCERNED
Possible candidates -- MARLENA JULIE DOUG ROMAN DON CHRIS KAYLA
Using MARLENA for WORRIER

>>> MARLENA is worried about RENEE

working on goal -- (TOGETHER * RENEE) -- using plan SEDUCTION

>>> RENEE seduces ROMAN

working on goal -- (ELIMINATE ALEX)

Several plans to choose from ATTEMPTED-MURDER EXPOSE
Using STEPHANO for KILLER

>>> STEPHANO tries to kill ALEX

...
```

Figure 10: Using a plot fragment with generalized constraints

4 Conclusion

There are many areas left to explore before our generalization methods can be fully implemented -- when to generalize a plot fragment, deciding exactly how much to generalize, and using a casual explanation to adjust details of the plot fragment, for example. However, we feel that this kind of explanation-based generalization of plot fragments, in conjunction with the basic story-telling methods we have developed and described here, will lead to dynamic systems that can generate wide ranges of interesting plot outlines. When coupled with the appropriate natural language generation techniques, such systems will produce interesting and exciting stories, as well as help us learn a great deal about many areas of cognitive processing, in particular, some of the more basic elements of creativity.

References

Appelt, D. E., 1982. Planning natural language utterances. Proceedings of the Second National Conference on Artificial Intelligence, Pittsburgh, PA, pp. 59 - 62.

Barthes, R., 1977. *Image -- Music -- Text.* Hill and Wang, New York.

Carbonell, J. G., 1983. Derivational analogy in problem solving and knowledge acquisition. Proceedings of the 1983 International Machine Learning Workshop, Champaign-Urbana, Illinois, pp. 12 - 18.

Charniak E., Riesbeck, C. K., and McDermott, D. V., 1980. *Artificial Intelligence Programming.* Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Dehn, N., 1981. Memory in story invention. Proceedings of the Third Annual Conference of the Cognitive Science Society, Berkeley, California, pp. 213 - 215.

DeJong, G. F., 1983. An approach to learning from observation. Proceedings of the 1983 International Machine Learning Workshop, Champaign-Urbana, Illinois, pp. 171 - 176.

Eco, U., 1979. *The Role of the Reader.* Indiana University Press, Bloomington, Indiana.

Eco, U., 1984. *Postscript to The Name of the Rose.* Harcourt Brace Jovanovich, New York.

Ellman, T., 1985. Generalizing logic circuit designs by analyzing proofs of correctness. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles.

Lebowitz, M., 1984. "Creating characters in a story-telling universe." *Poetics* 13:171 - 194.

Lehnert, W. G., 1981. "Plot units and narrative summarization." *Cognitive Science* 5:293 - 332.

McDonald, D. D., 1980. *Natural language generation as a process of decision making under constraint.* Ph.D. Thesis, MIT.

McDonald, D. D. and Conklin, E. J., 1982. Salience as a simplifying metaphor for natural language generation. Proceedings of the Second National Conference on Artificial Intelligence, Pittsburgh, PA, pp. 75 - 78.

McGuire, R., 1980. Political primaries and words of pain. Unpublished manuscript.

McKeown, K. R., 1982. *Generating natural language text in response to questions about database structure.* Ph.D. Thesis, University of Pennsylvania.

Meehan, J. R., 1976. The metanovel: Writing stories by computer. Technical Report 74, Yale University Department of Computer Science.

Michalski, R. S., Carbonell, J. G. and Mitchell, T. M., eds., 1986. *Machine Learning, An Artificial Intelligence Approach, Volume II.* Morgan Kaufmann, Los Altos, CA.

Mitchell, T. M., 1983. Learning and problem solving. Proceedings of the Eighth International Joint Conference on Artificial Intelligence, Karlsruhe, West Germany, pp. 1139 - 1151.

Mooney, M. and DeJong, G. F., 1985. Learning schemata for natural language processing. Proceedings of the Ninth International Joint Conference on Artificial Intelligence, Los Angeles, pp. 681 - 687.

Mostow, J., 1983. Operationalizing advice: A problem-solving model. Proceedings of the 1983 International Machine Learning Workshop, Champaign-Urbana, Illinois, pp. 110 - 116.

Sacerdoti, E., 1977. *A Structure for Plans and Behavior.* Elsevier North-Holland, Amsterdam.

Schank, R. C., 1972. "Conceptual Dependency: A theory of natural language understanding." *Cognitive Psychology* 3:532 - 631.

Schank, R. C., 1975. The structure of episodes in memory. In D. Bobrow and A. Collins, Ed., *Representation and Understanding: Studies in Cognitive Science*, Academic Press, New York .

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.* Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Todorov, T., 1977. *The Poetics of Prose.* Cornell University Press, Ithica, New York.

Turner, S. R and Dyer, M. G., 1985. Thematic knowledge, episodic memory and analogy in MINSTREL, a story invention system. Proceedings of the Seventh Annual Conference of the Cognitive Science Society, Irvine, CA, pp. 371 - 375.

Wilensky, R., 1983. *Planning and Understanding.* Addison-Wesley, Reading, MA.

Yazdani, M., 1983. Generating events in a fictional world of stories. Technical Report R-113, Computer Science Department, University of Exeter.