

Story Telling and Generalization

Michael Lebowitz

Department of Computer Science -- Columbia University
New York, NY 10027

Abstract

The generation of extended plots for melodramatic fiction is an interesting Artificial Intelligence task -- 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 might be automatically extended using explanation-based generalization methods. Our methods are based on analysis of a television melodrama.

1 Introduction

In [Lebowitz 84] and [Lebowitz 85] 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 AI problems. We introduced a story-telling program, UNIVERSE, concentrating on its ability to create realistic characters [Lebowitz 84] and a scheme for generating plot outlines using plan-like units, "plot fragments" [Lebowitz 85]. In this paper, we briefly describe our model of story telling, showing how UNIVERSE generates a simple piece of a plot outline. We then discuss how the appropriate way to extend the program is for it to automatically expand its knowledge base using techniques closely allied to explanation-based learning methods (e.g., [DeJong 83]). This discussion is based on analysis of plot outlines from a television melodrama.

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 reader/viewers. Eventually, we expect the program 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 low-level event generation for later. As an illustration of the kind of stories we would like to generate, consider the following synopsis of events from the 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 he was serious by a bomb exploding 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. So later, when Liz was finally free from Tony, Neil was not free to marry her and their troubles went on.

STORY1 exemplifies the kind of plot outline we wish 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. On the other hand, it is not enough that we simply simulate the lives of these characters (as in TALE-SPIN [Meehan 76]), since it is unlikely that they would naturally do such interesting things or that their actions would interleave with each other so nicely. Both our generation scheme and plan for generalizing new plot fragments 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 (although, since we are generating extended stories, the program must "tell as it plans"). It is based around a set of "plot fragments", which are initially built-in by the program designer, that play the same role as standard plans in planning systems such as [Sacerdoti 77; Wilensky 83]. We will discuss in Section 3 how the library of plot fragments might be extended automatically. The plot fragments provide narrative methods to achieve goals. What is crucial is that the goals and plot fragments are not viewed as goals and plans of the characters (although these must also be monitored), but, instead, goals and plans of the *author* (or program, in our case). This allows actions that make sense but yet are not necessarily what a character, if an independent agent, would choose to do. This approach fits nicely with work in narrative theory such as [Barthes 77; Eco 79; Todorov 77], that has influenced the development of UNIVERSE (as has the AI work of [Meehan 76; Dehn 81; Yazdani 83]).

The plot fragments used by UNIVERSE can span a wide range of levels -- from very general, thematic plans that may take a long time to carry out, to plans for specific actions. They may include steps with the sole purpose of setting the stage for later events (e.g., *dropping-hints*). The more general plot fragments are much like the plot units in [Lehnert 81], such as *double-crossing*. Other plot fragments are more specific, such as *seduction* or *wild-fight*, but the range of the fragments, along with the use of many characters, keeps the plot outlines interesting. Before starting a new story, UNIVERSE builds up a story-telling universe of characters, keeping track of personality traits, interpersonal relationships and goals for each, as described in [Lebowitz 84]. This proves to be precisely the way that some authors work in developing novels: create a set of characters and develop the plot from there (see [Eco 84]).

Figure 1 shows the main features of a typical UNIVERSE plot fragment, *forced-marriage*. It involves a nasty parent (?parent) forcing his daughter-in-law (?her) to stay in an unhappy marriage (to ?husband), preventing her from being with the person she really loves (?him), roughly the plot of STORY1.

The first piece of information about the *forced-marriage* plot fragment is the goal that it can be used to achieve -- "churn", keep two lovers from being happy. Obviously this goal makes no sense from the point of view of the characters, but makes a great deal of sense for the author, and, indeed, is a staple of melodrama. The next pieces of information in Figure 1 are a list of the roles involved in the plot fragment and constraints upon them. (Generalization of the constraints will be an important part of expanding the plot fragment

PLOT FRAGMENT: forced-marriage

```
GOALS: (churn ?him ?her) {prevent them from being happy}
TIME SCALE: months
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)
SUB-GOALS: (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 1: A typical UNIVERSE plot fragment

library.) For the roles that are not determined by the goal, UNIVERSE will find characters that fit the constraints (or create characters, if need be).

Finally, Figure 1 shows the heart of the plot fragment -- a series of subgoals (actually, a partial ordering) that must be achieved to attain the fragment's goal. Often, a plot fragment will include actual plot actions to be generated, but since *forced-marriage* is a relatively high level plot fragment, it simply spins off the series of sub-goals shown. Each of the subgoals in Figure 1 can potentially be satisfied by a variety of different plot fragments, leading to a wide range of possible stories.

The basic UNIVERSE story-telling algorithm is relatively simple, relying on the richness of its plot fragment library and character set. UNIVERSE maintains a precedence graph indicating the prerequisites of pending author goals and plot fragments. The program selects to pursue an author goal with all its prerequisites satisfied and a plot fragment to use for it (trying to achieve extra author goals, if possible). It continues this process as long as unsatisfied goals remain, generating concrete actions when appropriate. This algorithm, much like the one used in TALE-SPIN [Meehan 76] or micro-TALE-SPIN [Schank and Riesbeck 81], but using author goals, is summarized in Figure 2. This algorithm can lead to dead ends -- branches of disjunctive plans that cannot be achieved -- but, interestingly, as long as this does not happen too often, it is not a problem, being much like the "red herrings" often introduced in television melodrama.

One of the positive side effects of this algorithm is that as multiple goals are pursued, various "plot lines", i.e., the pursuit of different high level plot fragments, will become interleaved. This is important in the production of intricate plot outlines.

Figure 3 shows a brief run of UNIVERSE. In this example, we have given the program the top-level goal of "churning" Liz and Neil's relationship (e.g., keeping them apart) and a secondary goal of getting Renee and Neil together.

The program trace in Figure 3 (with plot outline output indicated by ">>>") shows the expansion of the *forced-marriage* plot fragment for the churn goal. Once UNIVERSE has

Liz becomes available). It picks the *seduction* plot fragment to do this. Rather than picking a seductress at random from among the characters with acceptable characteristics, UNIVERSE was able to satisfy its second goal as a side-effect by selecting Renee as the seductress. This sort of opportunistic planning goes a long way towards achieving the intricate interconnections that exist in most popular melodrama.

3 The next step: Generalizing plot fragments

We feel the generation framework outlined in Section 2 and described more fully in [Lebowitz 85], which is able to generate moderately interesting plot outlines, provides a good basis for future work. To generate interesting stories over the long run, 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 DA, who wanted to marry Hope for political purposes. Just as they were about to be married, Bo arrived, the danger past, and told Hope that he loved her. They ran off together. However, Maxwell (a new bad guy), 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 in fact married Larry, convincing Bo she no longer cared for him. Then, Megan, 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 4.

Figure 4 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 all these events play the same role -- they create a reason to keep the lovers apart. The I events, again, have rather different actions playing the same role -- to further thwart the lovers even when the prime obstacle between them disappears.

Notice that while all of these three plot lines are 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 1, or a somewhat more complex version of it). *forced-marriage* has rather specific constraints and calls for rather specific events. However, in STORY2B, instead of an overt threat from the woman's

Pick a goal with no missing pre-conditions

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

**“Execute” the plot fragment
(generating plot actions and adding new goals to the goal graph)**

Repeat

Figure 2: The basic UNIVERSE story-telling algorithm

```
*(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-SWEAK-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 3: Generating a simple plot outline

selected this plot fragment, each of the subgoals is pursued in turn. As it happens, in this example each subgoal can be achieved by a single plot fragment. In a larger example, of course, some of these subgoals would lead to further expansion of subgoals, which might intertwine with each other.

An interesting point in Figure 3 comes when UNIVERSE tries to satisfy the (together * neil) goal, e.g., get Neil involved with someone (so that he will be entangled romantically when

| <u>STORY1</u> | <u>STORY2A</u> | <u>STORY2B</u> |
|--|---|--|
| 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 4: Three similar plot outlines

father-in-law there is a covert threat from the danger Bo is in. We could, however, hope to derive from *forced-marriage* a more general plot fragment that would produce STORY2 .

Based on this sort of analysis of existing melodramatic plot outlines, we propose to automatically augment UNIVERSE's library of plot fragments by generalizing existing plot fragments.¹ 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 event B involved a competing goal, step D a threat (optional), step G a competing relationship and so forth.

¹An 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. Generalizing new plot fragments fits better into the UNIVERSE scheme.

Then this plot fragment could be instantiated into either STORY2A or STORY2B (or back into STORY1).

The main process of generalization 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 within the generalization. So, for example, instead of requiring a married couple in *forced-marriage*, we might have UNIVERSE just look for any sort of a normally positive relationship. Various details will be added to the actions and sub-goals depending on how the plot fragment is instantiated.

While it is easy to simply generalize the role constraints (in fact, 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 and 2) making sure they 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 83].² The basic idea is that we would build up a causal analysis of a plot fragment, and then generalize it in various ways such that the analysis still holds, as done in [DeJong 83] to build up new story understanding schemata. Using the causal analysis, built up in ways such as those described in [Schank and Abelson 77], 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 83]. Both of these differences address the second of our concerns -- that the generalized plot fragments still yield interesting stories. We propose 1) that author goals be included in the explanations of the plot fragments and 2) that we not generalize plot fragments 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. 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

²There has been considerable other work into explanation-based learning (see [Michalski et al. 83]). However, DeJong's story understanding domain is most closely related to our work.

plot fragment is unlikely to produce interesting stories, as it will not further author goals that revolve around intense relationships.

Our second point is that we do not wish to generalize plot fragments as far as possible. 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, incidentally, can also make use of the causal analysis needed for generalization).

Even small generalizations of the 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.

As a partial illustration of generalizing plot fragments, we have developed a simple module to generalize constraints without regard to underlying causality (which must, of course, be considered eventually). Figure 5 shows the generalized constraints for *forced-marriage*. We can see 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 5: Generalized constraints for *forced-marriage*

Figure 6 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 3. 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-MARRIAGEO {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 6: 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, 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.

References

- [Barthes 77] Barthes, R. *Image -- Music -- Text*. Hill and Wang, New York, 1977.
- [Dehn 81] Dehn, N. Memory in story invention. Proceedings of the Third Annual Conference of the Cognitive Science Society, Berkeley, California, 1981, pp. 213 - 215.
- [DeJong 83] DeJong, G. F. An approach to learning from observation. Proceedings of the 1983 International Machine Learning Workshop, Champaign-Urbana, Illinois, 1983, pp. 171 - 176.
- [Eco 79] Eco, U. *The Role of the Reader*. Indiana University Press, Bloomington, Indiana, 1979.
- [Eco 84] Eco, U. *Postscript to The Name of the Rose*. Harcourt Brace Jovanovich, New York, 1984.
- [Lebowitz 84] Lebowitz, M. "Creating characters in a story-telling universe." *Poetics 13*, 1984, pp. 171 - 194.
- [Lebowitz 85] Lebowitz, M. Generating plot outlines. Columbia University Department of Computer Science, 1985.
- [Lehnert 81] Lehnert, W. G. "Plot units and narrative summarization." *Cognitive Science 5*, 4, 1981, pp. 293 - 332.
- [Meehan 76] Meehan, J. R. The metanovel: Writing stories by computer. Technical Report 74, Yale University Department of Computer Science, 1976.
- [Michalski et al. 83] Michalski, R. S., Carbonell, J. G. and Mitchell, T. M., eds. *Proceedings of the 1983 Machine Learning Workshop*. University of Illinois, Urbana-Champaign, IL, 1983.
- [Sacredoti 77] Sacredoti, E. *A Structure for Plans and Behavior*. Elsevier North-Holland, Amsterdam, 1977.
- [Schank and Abelson 77] Schank, R. C. and Abelson, R. P. *Scripts, Plans, Goals and Understanding*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1977.
- [Schank and Riesbeck 81] Schank, R. C. and Riesbeck, C. K. *Inside Computer Understanding*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1981.
- [Todorov 77] Todorov, T. *The Poetics of Prose*. Cornell University Press, Ithica, New York, 1977.
- [Wilensky 83] Wilensky, R. *Planning and Understanding*. Addison-Wesley, Reading, MA, 1983.
- [Yazdani 83] Yazdani, M. Generating events in a fictional world of stories. Technical Report R-113, Computer Science Department, University of Exeter, 1983.