

Planning Stories

Michael Lebowitz¹

Department of Computer Science -- Columbia University

New York, NY 10027

(212) 280-8196

4 February 1987

Full paper; Keywords: Story generation; planning

Abstract

Story generation can best be viewed as a planning task. We show here UNIVERSE, a program that generates melodrama plot outlines using hierarchical planning methods. Examples are given of the program creating story outlines using a set of characters that it also created. We indicate that story telling is open-ended, does not have to be perfect, and evaluation criteria are unclear, and contrast the sort of planning needed for story telling with other planning tasks. We suggest that certain elements of the methods used by UNIVERSE could be usefully applied to other tasks.

1 Introduction

The generation of stories is a challenging problem for Artificial Intelligence techniques. We have designed a program, UNIVERSE, that creates story outlines in the domain of interpersonal melodrama, a common form of which is soap opera. We chose this domain because it revolves around character relations, rather than action. It allows us to look at cognitive science issues such as author intention, knowledge-state assessment, character representation and eventually a range of natural language issues, in a very accessible domain. In this regard we have studied how narrative theory work such as (Barthes, 1977; Eco, 1979) applies to our task. The domain also has long-run potential in the areas of the education and interactive entertainment.

With UNIVERSE we view story telling primarily as a planning process, much in the same way that (Cohen and Perrault, 1979; Appelt, 1985) view language generation. This contrasts with, for example, making direct use of story grammars (Rumelhart, 1975; Mandler and Johnson, 1977) to generate a story.² In this paper we will briefly describe how UNIVERSE plans a story outline and then discuss some of the more interesting ways in which this sort of planning contrasts with planning in more traditional AI settings. Further details of UNIVERSE can be found in (Lebowitz, 1984; Lebowitz, 1985).

We have designed UNIVERSE to create plot outlines -- the major events that happen in a story -- since we are not yet ready to deal with the problems inherent in dialogue. These outlines are much like the summaries of soap operas that often appear in newspapers, or, these days, on computer services. EX1 shows a CompuServe[®] summary for a television melodrama. This is typical of the kinds of outlines that we have UNIVERSE produce, although UNIVERSE usually deals with broader events over a longer period of time.

¹A series of projects by Paula Langer and Doron Shalmon contributed greatly to UNIVERSE.

²See (Black and Wilensky, 1979; Mandler and Johnson, 1980) for further discussion of story grammars.

EX1 -

Days of Our Lives-12/26/86

At her wit's end, Kim begs Shane to use the new ISA truth serum on Emma. They administer the drug and start questioning her but it is unclear as to whether Emma is really faking it or not. Poor Carrie is upset when Roman leaves for work. She is afraid that he won't be coming back. Frankie comforts her and they decide to watch the VCR together. When Frankie leaves to rent some movies, he finds her gone when he returns. He guesses that Carrie ran to the park to be with Roman. In the park, the criminal grabs her but Frankie bravely rescues her. Roman assures Carrie that he can handle his work. Finding Kayla without an assistant at the clinic, Patch offers to help out. He helps sign the patients in and keeps their spirits up much to Kayla's delight. While there, the mystery woman keeps watch on Patch. Later, Patch catches her and demands to know why she is following him.

We can see from EX1 that interpersonal melodrama involves a wide range of events and characters. The plots are quite intricate and interwoven, playing off of the personalities of the characters. The plot interconnections of EX1 are clearer if you know that Kim, Kayla and Roman are siblings, Shane is Roman's ex-partner, and Patch and Roman have a long-standing feud.

2 Story telling as planning

UNIVERSE uses a hierarchical planning algorithm much like that of NOAH (Sacerdoti, 1977). It knows about a variety of goals and uses plot fragments to achieve them. Much like NOAH's plans, or those used for story understanding in PAM (Wilensky, 1983), the plot fragments consist of a series of subgoals and have associated with them a set of constraints, mostly involving the kinds of characters that can be used as role fillers, that determine when they can be used. Currently UNIVERSE has about 65 plot fragments in its library, some of which will be seen below. This is a large enough set to convince us of the validity of our methods, and generation some interesting plot outlines, although clearly a much larger set would be needed for a full-fledged program.

The key to UNIVERSE's planning algorithm is that it does not concentrate on the goals of the characters, but instead on *author* goals. In this way the planning process leads to stories that are interesting and have points. If we simply use character goals, then the characters will behave believably, but probably not interestingly. To illustrate this point, UNIVERSE uses one common goal *churn*, to keep apart two characters who are in love (by putting obstacles in their way). Clearly this is not a goal of the characters, but it leads to more interesting stories than simply having them live happily ever after. The use of author goals is the key distinction between UNIVERSE and Meehan's TALE-SPIN (Meehan, 1976) which in many ways was the inspiration for our program. TALE-SPIN would simulate characters responses to goals in Aesop's fables situations. MINSTREL (Turner and Dyer, 1985), a more recent story-telling program, does consider author goals to a point, but is more concerned with memory issues than we are here, and is still basically character oriented.

In developing UNIVERSE, we rapidly came to the conclusion that planning would be simplified if we created a set, or universe, of characters before beginning to tell stories (hence the name of the program). The rationale behind this is that the constraints needed to create believable characters -- e.g., creating all the marriages and divorces that add "color" to the story -- were unlike those for the rest of the story-telling process. UNIVERSE can create characters for a plot fragment "on the fly" when none of the

Lebowitz: Planning Stories

existing ones are appropriate, but always doing so would decrease the coherence of the story. Each character is represented by a frame that describes parents, spouses, descriptive stereotypes (for coherence) past historical events, and, most importantly, various character traits and interpersonal relations. These values are used as constraints on the various plot fragments that can be applied. (Lebowitz, 1984) describes the character representation and creation process in detail. Figure 1 provides English summaries of a few of the characters created by UNIVERSE that will be used in the examples below.

```
Fran      -- a 25-year-old, nice, able, bureaucrat
Joshua    -- Fran's husband, a sleazy New York lawyer who is surprisingly nice
Valerie   -- another, not-so-nice, sleazy lawyer
Louis     -- Valerie's preppie husband
Gerald    -- Joshua's father, a nice fellow, who, nonetheless, recently
           had an affair with Fran
```

Figure 1: A partial cast of characters

Given a set of plot fragments, a set of characters, and a "seed goal", UNIVERSE operates much like NOAH (which turns out to be fairly similar to TALE-SPIN and even more so to micro-TALE-SPIN (Charniak et al., 1980)). It maintains a precedence graph with a partial ordering of which goals must be achieved before others. It uses a least commitment, opportunistic planning algorithm that repeatedly:

- Selects a goal with no unfulfilled goals that must precede it.
- Finds the relevant plot fragments that should achieve the goal, along with possible role bindings.
- Picks one of these plot fragments.
- Expands the selected plot fragment; this may include: 1) creating new goals, 2) modifying interpersonal relations and/or character traits, and 3) generating text (which is currently done from simple templates).

This simple algorithm appears to generate reasonable plot outlines even though our plot fragment library is still relatively small. We compare it to more traditional planning in Section 3. A few points about the algorithm are worth noting. UNIVERSE does not completely plan out a story before it generating any of it. This is because it is intended for open-ended story situations that may have no natural ending. For example, most of the plot fragments for **churn** have as their final subgoal to **churn** further. Generating as it goes leaves UNIVERSE open to the problem of running into blind alleys where it has already generated part of a story that it cannot complete. We discuss this further below, but simply note here that the program can often re-plan from the state achieved in the blind alley.

An important part of the UNIVERSE algorithm is deciding upon the plot fragment to choose when many will achieve the author goal. Since we believe in the necessity of a large plot fragment database, this is a very real issue. Indexing allows UNIVERSE to easily find the relevant fragments, but we still need a way to select one. For the moment, UNIVERSE uses two selection criteria. First it checks for plot fragments that will achieve other open goals in the precedence graph. This method, which will be illustrated below, leads to nice connections among various plot threads. The second criterion is to use interest values associated with each fragment. We feel that these interest levels should actually also depend on the characters involved (Lebowitz, 1981). If there are still multiple possible fragments, UNIVERSE selects one randomly.

Lebowitz: Planning Stories

Figure 2, which shows a fairly mundane beginning to a plot outline planned by UNIVERSE using characters that it created, will be used to illustrate the algorithm. Lines preceded by ">>>" are the plot outline output.

```
*(tell '(((churn JOSHUA FRAN)))

working on goal -- CHURN JOSHUA FRAN
Several possible plans with equal goal effect
LOVERS-FIGHT HIM/FRAN HER/JOSHUA
LOVERS-FIGHT HIM/JOSHUA HER/FRAN
JOB-PROBLEM P1/FRAN P2/JOSHUA JOB/BUREAUCRAT
JOB-PROBLEM P1/JOSHUA P2/FRAN JOB/SLEAZY-LAWYER
PREGNANT-AFFAIR WOMAN/FRAN HUSBAND/JOSHUA
ACCIDENT-BREAKUP P1/FRAN P2/JOSHUA
STEAL-CHILD HUSBAND/JOSHUA HER/FRAN
COLLEAGUE-AFFAIR WIFE/FRAN HUSBAND/JOSHUA
AVALANCHE-ACCIDENT HIM/FRAN HER/JOSHUA
AVALANCHE-ACCIDENT HIM/JOSHUA HER/FRAN
  -- picking one
  -- using plan LOVERS-FIGHT HIM/FRAN HER/JOSHUA

working on goal -- DO-FIGHT JOSHUA FRAN
Several possible plans with equal goal effect
PERSONAL-FIGHT A/JOSHUA B/FRAN SUBJECT/IN-LAWS
PERSONAL-FIGHT A/JOSHUA B/FRAN SUBJECT/MONEY
PERSONAL-FIGHT A/JOSHUA B/FRAN SUBJECT/SECRETS
PERSONAL-FIGHT A/JOSHUA B/FRAN SUBJECT/FLIRTING
PERSONAL-FIGHT A/JOSHUA B/FRAN SUBJECT/KIDS
  -- picking one
  -- using plan PERSONAL-FIGHT A/JOSHUA B/FRAN SUBJECT/MONEY

>>> JOSHUA and FRAN fight about: MONEY
working on goal -- DUMP-LOVER JOSHUA FRAN
  -- using plan BREAK-UP DUMPER/JOSHUA DUMPED/FRAN

>>> JOSHUA tells FRAN he doesn't love her

working on goal -- CHURN FRAN JOSHUA
Several possible plans with equal goal effect
LOVERS-FIGHT HIM/JOSHUA HER/FRAN
JOB-PROBLEM P1/JOSHUA P2/FRAN JOB/SLEAZY-LAWYER
JOB-PROBLEM P1/FRAN P2/JOSHUA JOB/BUREAUCRAT
PREGNANT-AFFAIR WOMAN/FRAN HUSBAND/JOSHUA
ACCIDENT-BREAKUP P1/FRAN P2/JOSHUA
STEAL-CHILD HUSBAND/JOSHUA HER/FRAN
COLLEAGUE-AFFAIR WIFE/FRAN HUSBAND/JOSHUA
AVALANCHE-ACCIDENT HIM/JOSHUA HER/FRAN
AVALANCHE-ACCIDENT HIM/FRAN HER/JOSHUA
  -- picking one
Multiple deferred filler -- picking one
  -- using plan COLLEAGUE-AFFAIR WIFE/FRAN HUSBAND/JOSHUA COLLEAGUE/JACK

[and the story continues with the protagonists living unhappily ever after']
```

Figure 2: The beginning of a UNIVERSE plot outline

The first action shown in Figure 2 is UNIVERSE, given the goal of churning two characters, Joshua and Fran, collecting all of the plot fragments that might satisfy that goal, filtering out those with constraints that are not met by Joshua and Fran. Since the churn author goal has been one of our prime examples, there are a number of possibilities. For each goal, UNIVERSE collects all the possible role bindings so that it can sensibly select among them.³ After selecting the lovers fight fragment, essentially at random, since there are no other goals and all the fragments have the same interest value, UNIVERSE expands its subgoals. For one of these, the fight itself, we can see that UNIVERSE considered a number of possible fight topics. The topics are generated from knowledge of the characters involved and so would be different for different characters.

After UNIVERSE finishes with the lovers fight, it decides to churn the relationship further. The selection of the colleague affair fragment illustrates an important point. After its selection, UNIVERSE has to fill more roles. We discovered that there were too many combinations to allow UNIVERSE to expand out all possibilities for the various minor characters in all the relevant fragments in order to pick the best. So, instead, each plot fragment has identified the roles that we are sure it can fill with somebody (by creating a new character if need be) and are probably not relevant to the author goals of the plot, and hence are not likely to be relevant to picking the best fragment to use. The filling of these roles is deferred until after a fragment is chosen, a technique that might be applicable to other planning situations.

Figure 3 shows a second example outline that illustrates how extra goals affect the opportunistic planning of UNIVERSE. Here, as well as seeding UNIVERSE with a churn goal, we have also asked it to get Joshua **together** with Valerie. This goal, while given the program here, could equally well have come from its own pursuit of other goals. The notable point is that now, when looking for a fragment for churn, it selects accident-breakup, the only one that guarantees that the man will get together with a desired person, given appropriate bindings. We show the outline as far as where UNIVERSE begins to plan the new relationship which will, in the long run, create interest by making Joshua unavailable should Fran get better. Opportunistic planning, does an excellent job in pulling story strands together. The ability to satisfy several goals with one plan seems to be a hallmark of clever melodrama.

3 Story planning versus action planning

The planning of stories done by UNIVERSE shares much with other sorts of AI planning methods. As mentioned above, the algorithm is very similar to various hierarchical planners in its use of goals and plans (and could probably be made more similar to good effect). However, it is interesting to contrast both the task and the method with standard AI action planning. We will also indicate how it may be advantageous to apply many aspects of story planning to more standard action planning.

The key points that we see in the planning of stories, each of which we will discuss below, are:

- You can “say as you go”; it is neither necessary nor possible to fully plan out a story before beginning to tell it.
- The plan does not have to be perfect; people will accept odd combinations (and indeed fill in motivation).

³Some of the role names are a bit deceptive. For example, the *him* and *her* roles of the lovers fight plot fragment are actually genderless, which is why UNIVERSE tries both bindings.

Lebowitz: Planning Stories

```
*(tell '(((churn JOSHUA FRAN)) ((together JOSHUA VALERIE))))  
  
working on goal -- CHURN JOSHUA FRAN  
-- using plan ACCIDENT-BREAKUP P1/FRAN P2/JOSHUA THIRD-PARTY/VALERIE  
  
working on goal -- DO-DISABLE FRAN  
-- using plan DISABLE PERSON/FRAN  
  
>>> FRAN has a spinal injury and is paralyzed  
  
>>> FRAN doesn't want to ruin JOSHUA's life  
  
>>> FRAN pretends to blame JOSHUA for her malady  
  
working on goal -- DUMP-LOVER FRAN JOSHUA  
-- using plan BREAK-UP DUMPER/FRAN DUMPED/JOSHUA  
  
>>> FRAN tells JOSHUA she doesn't love him  
  
working on goal -- TOGETHER JOSHUA VALERIE  
  
[again, the story continues unhappily for almost all concerned]
```

Figure 3: A multi-goal story

- Selecting the plot (plan) to use is important; heuristics that consider factors other than goal satisfaction must be used.
- Unlike most action domains, each role filler (character) is different, and they cannot be used interchangeably.

Many forms of stories, including soap opera and some children's stories, have no real end. As a result, we clearly cannot plan such stories in their entirety before generating any of the story. At the moment, UNIVERSE probably goes too far in simply generating text that appears to lead to the goal, often getting into binds that it cannot resolve. On the other hand, as we will mention below, this is not that crucial for a story. In examining television serials, one can frequently see situations where the writers have clearly headed down blind alleys of this kind. An important future research topic will be to determine just how far one should plan a story -- short of guaranteeing goal satisfaction, which is not practical -- before actually generating any of it.

We feel that more planning situations would be appropriate for an "execute as you go" paradigm than are currently approached that way. For example, if we are developing a program to plan a series of financial investments, we may want it to start and make some of the moves before waiting to plan everything. Since much of the plan is likely to be dependent on what happens at the beginning, there is no real point in planning all the way to the end. Admittedly, done more completely than by UNIVERSE this sort of planning is quite complicated. At a minimum, it involves deciding when you've planned enough, feedback from execution, and re-planning, but should be well worthwhile.

One factor that makes "say as you go" particularly feasible for story telling is that story outlines do not have to be perfect. They can go into blind alleys and backtrack in unusual ways and still be accepted by readers. This is true for at least two reasons. Obviously, there is no such thing as a single *best* story. As long as a story is coherent and does not blatantly violate rules of the world, then it is, in some sense, acceptable. Indeed, it is unlikely that at present a program could detect fine differences in story quality. A secondary reason for the lack of need to look for the best story is that our human readers will often fill in

details that the program may not have had in mind (Schank and Abelson, 1977). Again, many domains beyond story telling also do not need optimal plans. For example, a program designed to come up with recipes should not try to evaluate minor differences in quality.

The next interesting point regarding story planning is that for many author goals there are typically a sizable number of different plot fragments that will achieve them, not just one or two. Further, we have seen that the program often cannot really evaluate the stories that result from different choices. Nonetheless, it is clear that some decisions will work out better than others. As a result, it is necessary that UNIVERSE pay considerable attention to which plot fragment to select, *even among those that all achieve the author goal*. In particular, it is necessary that the program use heuristics that consider factors other than simple goal satisfaction. We feel that local factors such as interest levels and patterns of character involvement can, along with the opportunistic planning that we have described, lead to intricate and interesting plot structure.

The final area of interest involving planning stories is that, unlike most action domains, every character in a fictional universe is different. In most action domains that have been studied, while there may be a number of objects in the world, they tend to fall into a few small classes (e.g., rectangular blocks) that can be treated identically by operators. In story telling, though, we have to be careful to make sure that the role fillers (characters) used by a plan are appropriate in order to retain believability. Of course, we also have the option of creating new characters out of nowhere which is not open to other forms of planning. While the diversity of role fillers does not apply to most domains, it may be relevant in situations where we are trying to understand stories involving people or human actions.

4 Conclusion

Story telling is a form of generation that involves a number of quite interesting AI problems. Planning in terms of author goals definitely seems to be the way to think about story telling, rather than any form of story grammar. We feel that the kind of planning done for story telling reflects what people do in many different planning situations (not just stories) and could be profitably applied to AI systems dealing with planning problems. In the future, we plan to expand the basically simple methods of UNIVERSE to make use of more complex planning methods to get further intricacy into our plot outlines, including, perhaps, case-based planning of the sort described in (Kolodner, Simpson and Sycara-Cyranski, 1985; Hammond, 1986), or thematic issues such as those considered by (Turner and Dyer, 1985).

References

- Appelt, D. E. (1985). *Planning Natural Language Utterances*. Cambridge, England: Cambridge University Press.
- Barthes, R. (1977). *Image -- Music -- Text*. New York: Hill and Wang.
- Black, J. B. and Wilensky, R. (1979). An evaluation of story grammars. *Cognitive Science*, 3, 213 - 229.
- Charniak E., Riesbeck, C. K., and McDermott, D. V. (1980). *Artificial Intelligence Programming*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Cohen, P. R. and Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. *Cognitive Science*, 3, 177 - 212.
- Eco, U. (1979). *The Role of the Reader*. Bloomington, Indiana: Indiana University Press.
- Hammond, K. (1986). CHEF: A model of case-based planning. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence* (pp. 267 - 271), Philadelphia, PA: Morgan Kaufmann.
- Kolodner, J. L., Simpson, R. L. and Sycara-Cyranski, K. (1985). A process model of case-based reasoning in problem solving. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 284 - 290), Los Angeles: Morgan Kaufmann.
- Lebowitz, M. (1981). Cancelled due to lack of interest. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence* (pp. 13 - 15), Vancouver, Canada: Morgan Kaufmann.
- Lebowitz, M. (1984). Creating characters in a story-telling universe. *Poetics*, 13, 171 - 194.
- Lebowitz, M. (1985). Story-telling as planning and learning. *Poetics*, 14, 483 - 502.
- Mandler, J. M. and Johnson, N. S. (1977). Remembrance of things parsed: Story structure and recall. *Cognitive Psychology*, 9, 111 - 191.
- Mandler, J. M. and Johnson, N. S. (1980). On throwing out the baby with the bathwater: A reply to Black and Wilensky's evaluation of story grammars. *Cognitive Science*, 4, 305 - 312.
- Meehan, J. R. (1976). *The metanovel: Writing stories by computer*. (Tech. Rep. No. 74). Yale University Department of Computer Science.
- Rumelhart, D. E. (1975). Notes on a schema for stories. In D. Bobrow and A. Collins, (Ed.), *Representation and Understanding: Studies in Cognitive Science* New York: Academic Press.
- Sacerdoti, E. (1977). *A Structure for Plans and Behavior*. Amsterdam: Elsevier North-Holland.

Lebowitz: Planning Stories

Schank, R. C. and Abelson, R. P. (1977). *Scripts, Plans, Goals and Understanding*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

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* (pp. 371 - 375), Irvine, CA.

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