A User Model for the Generation of Dilemma-based Interactive Narratives

Heather Barber and Daniel Kudenko
University of York
Heslington, York, YO10 5DD
email: {hmbarber,kudenko}@cs.york.ac.uk

Abstract

This paper investigates a user model for a system which automatically generates interactive stories. These stories are focused on dilemmas in order to create dramatic tension. The system is provided with knowledge of generic story actions and dilemmas based on those clichés encountered in many of today's soap operas. These dilemmas and story actions are instantiated for the given storyworld and a story planner creates sequences of actions that each lead to a dilemma for a character (who can be the user). The user interacts with the story by making decisions on relevant dilemmas and by freely choosing their own actions. In this paper we propose a method of generating a model to predict user decisions. This model is used to optimise the user's experience - by selecting the next dilemma to be presented according to what will be the most interesting and by looking ahead when planning.

Introduction

In recent years computer games from most genres have included a progressive story line to increase the immersive experience of the user and their enjoyment of the game. However, such stories are often linear, and in almost all cases pre-defined, which reduces the replay value of these games. Research into interactive narrative generation (or interactive drama) tries to overcome these weaknesses. Most interactive drama systems (prominent examples include (Young 2001; Bates 1992; Cavazza & Charles 2002; Crawford 2004; Fairclough 2004; Rousseau & Hayes-Roth 1998; Karlsson *et al.* 2006; Magerko 2005; Mateas & Stern 2003; Sgouros 1997; Szilas 2003; Thomas & Young 2006)) are focused on generating short story lines and do not adapt to the user (see Section "Related Work" for exceptions).

In this paper, we propose a system that generates interactive stories which are long (potentially infinitely so), and that adapt to the user's behaviour. This paper focuses particularly on the user model. To add dramatic tension, the story incorporates dilemmas as decision points for the user. These dilemmas are based on the clichés found in many contemporary soap operas, such as the trade-off between personal gain and loyalty to a friend. Overarcing stories connect these dilemmas as points of dramatic tension within a coherent plotline which is dynamically created, based on the user's response and action choices.

Our goal is to keep the story designer's input to a minimum and the user involvement as high as possible. In the proposed system, the story designer provides the story background in the form of character information and other knowledge that relates to the world in which the story is to be created (e.g. the east end of London). The system then instantiates all generic knowledge on story actions and dilemmas accordingly and thus creates the narrative in collaboration with the user's actions.

When presented with a dilemma the user is free to make a decision. This updates the user model, so the system can predict likely future user decisions. This enables it to select future dilemmas according to which will lead to the most interesting outcome. Here we present an early design for the user model which has shown promise in early testing.

This paper is structured as follows. First a general overview of the system is given, followed by a discussion of the story background representation, a description of dilemmas and the story generator. We proceed with a discussion of the user model motivation, design, testing and extensions. The paper finishes with a brief overview of related work and conclusions.

System Overview

The proposed knowledge-based interactive drama system is introduced here. Figure 1 shows the overall structure and the interactions between the system components. The interactive drama knowledge base consists of: the storyworld (which contains information regarding the characters); story actions; and dilemmas which can occur in the storyworld. This information is partially genre dependent and provided by the story designer, with the remainder being hard coded. These components are drawn upon in the generation of a narrative through planning. The user is able to interact with the narrative generator, and their actions effect the story experienced. A user model is employed to ensure that the story's dramatic interest is maximised.

The Storyworld

The storyworld consists of characters and locations at which the characters can be. These characters have various associ-

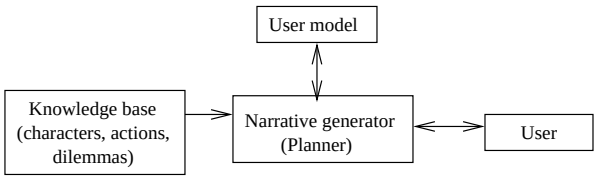

Figure 1: This figure shows the components of the system and how they interact.

ated attributes (e.g. attractiveness, gender) and characteristics (e.g. generosity and morality) which each have a range of values associated. Characters have varying strengths of unidirectional relationships with one another. There are a range of principles for which each character has a strength of belief – this can be reduced under certain pressures and circumstances. Characters also have aspirations, for example wanting a baby.

A character's nature, principles and aspirations affect which actions and dilemmas they can participate in and also, ideally, the user's opinion of that character. Each character should act in a manner which is consistent with their traits and how they have acted previously, while at the same time avoiding predictability.

A series of genre-specific locations are required by the storyworld. At any given time in the story, each character is at one of these locations. Direct interactions between characters can only take place if they are at the same location.

Actions

Those actions which can take place within the storyworld must be specified for each domain. Every possible action should be included and although these vary between domains there remains a significant overlap. These can include characters falling in love, becoming pregnant and being involved in crimes – such as drugging or murder.

Each action has associated conditions which must be satisfied before execution (preconditions) and effects representing changes to the storyworld following execution. For example, the action of a character moving between locations l and k has preconditions of the character being at location l and the existence of a path between locations l and k. The effects of this action are that the character is at location k and is not at location l. This follows the STRIPS representation.

Before an action is made available to the system for use within a storyline an applicability check is carried out. An action can only be utilised if its applicability is high enough. This ensures that the action is of the type that the acting character is likely to make. For example, a more attractive character starting to fancy a very generous character.

Every act that other characters within the system can make is available to the user who is able to freely specify their own actions within the scope of the current genre. The user inputs their action choices as two or three typed words which summarise the action they have chosen, for example 'move club' to move from their current location to the club. The system recognises a range of possibilities for each action. Additional options available to the user include being able to see the current state of the storyworld.

Dilemmas

Field (Field 1984) states that "drama is conflict", that the dramatic interest in a story centralises on its conflicts. In genres which make use of clichéd storylines these are usually found to be essentially conflicts (or dilemmas). Writers utilise these dilemmas in the creation of stories. A general form of each such clichéd dilemma can be determined, and a computerised story writer can create an interactive drama around these. Dilemmas require characters to make fundamentally difficult decisions within the course of the story.

Our experience showed that when more than two characters were involved in a dilemma it was either expandable to multiple two character dilemmas, or the characters receiving payoffs naturally divided into two groups with the same resultant utility. Therefore a decision on a dilemma involves only two recipients of utility payoffs. Five such dilemma categories were identified. These consist of all situations with two payoff recipients where there is a dilemma involved. This may require characters to be friends or enemies. The relevant categories are: Betrayal, Sacrifice, Greater Good, Take Down and Favour. Further details are given on each dilemma type in the following subsections.

Betrayal

When presented with a Betrayal dilemma a character must decide whether or not to take an action which would result in their best possible utility but simultaneously the worst possible outcome for their friend (or someone close to them). The decision would clearly not involve a dilemma were the two characters not friends. A character having the opportunity to be unfaithful to their partner is an example of the Betrayal dilemma.

Sacrifice

A character facing the Sacrifice dilemma is able to choose an action which will result in their worst possible utility but also the best possible outcome for their friend. It is necessary that these characters are friends for this to be a dilemma. An example of the Sacrifice dilemma occurs when a character has committed a crime which their friend has been accused of. Here the character has the opportunity to admit to their crime and thus accept the punishment rather than allowing their friend to take the blame.

Greater Good

Involvement in a Greater Good dilemma means that a character is able to take an action which will result in their best possible utility but also the best outcome for their enemy. This would not be a dilemma if the characters were not enemies. An instance of the Greater Good dilemma involves a character deciding whether to give something (such as information or a friend) to their enemy in order to save themself.

Take Down

In the Take Down dilemma a character has the option of taking an action which will result in their worst possible utility but also the worst outcome for their enemy. Clearly the characters must be enemies for the dilemma to exist. A character deciding whether to injure (or even kill) their enemy in full awareness that they will receive a punishment for this crime would be involved in the Take Down dilemma.

Favour

The favour dilemma sees a character X able to choose whether or not to take an action where there will not be any immediate discernible benefit to X as a result of their decision. The utilites of characters Y and Z will change as a result of this action decision. If X chooses to take the action the outcome will be the best possible for Y and Z will receive their lowest utility – and vice versa if X chooses not to take this action. An instance of this dilemma occurs when a character must choose between potential partners.

As can be seen, the Betrayal and Sacrifice dilemmas are the inverse of one another, as are the Greater Good and Take Down dilemmas. This means that any dilemma which falls into one of these categories can be inverted to become a dilemma of the other category. All five categories are kept to increase ease of dilemma identification within specific genres. From these categories dilemma instances can be found and generalised for each domain. From the generalised form of the dilemma the system will be able to create new dilemmas. In the presentation of these wholly original stories are created.

It will not be possible to create great literature in this way – the use of clichéd storylines prevents this. However, such stories are enjoyed by many people and this method is common in such genres as James Bond films, soap operas (soaps) and "chick flicks". The story is built around the cliché, and it is the cliché as well as the story which the audience appreciate, the very repetitiveness and familiarity of the dilemmas adding to the dramatic interest. Much more enjoyment could arise from the user becoming a character in such domains, and experiencing the dilemmas.

The Narrative Generator

Prior to a dilemma being presented certain conditions must be met within the storyworld. These are the preconditions of the dilemma. It is the task of the storywriting system to achieve these preconditions. Given actions (including those for the user) within the storyworld the system can plan to satisfy a dilemma's preconditions. A plan to achieve a dilemma thus becomes a storyline. The interactive drama is made up of a series of such substories, dynamically selected according to appropriateness.

On being passed a dilemma, the planner finds all plans to achieve this dilemma given the current storyworld state and background knowledge. From these plans, that which is most dramatically interesting can be selected and execution attempted. If the plan is successful the corresponding dilemma is presented. Once a decision has been made the system updates the storyworld state accordingly. The system can then plan from the new state for another dilemma – thus continuing the interactive drama. This sequence of events is demonstrated in fig. 2.

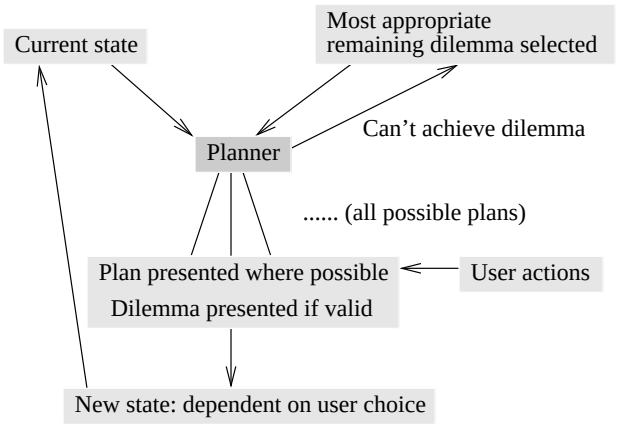

Figure 2: This figure gives an overview of the system moving between states dependent on plans, dilemmas and user decisions.

The sequence in which the dilemmas are selected for planning must depend on what has happened previously to become part of a consistent story. The frequency of dilemma use will need to be determined for each dilemma in all domains and considered when selecting each dilemma. It must be ensured that the user experiences a reasonable proportion and balance of dilemmas while the overall frequency is as would be expected for the genre. Dilemmas with greater dramatic interest are preferred.

All dilemmas are possible for any characters – including the user – within the storyworld (given applicability and satisfaction of preconditions). This increases the user's belief in the characters. Their decision making will mean they seem more realistic, as their character traits will become clear through this.

The planner initially assumes that the user will act in a manner consistent with the way characters with similar traits act in soaps. Once a plan has been chosen its actions are presented until the preconditions of an action or the dilemma cannot be satisfied without the user's participation. If the user acts in a manner which satisfies the necessary preconditions at this stage then the presentation of the plan continues until a user action is required again. As soon as it becomes possible to present the dilemma this is done. The user can be coerced into acting in the required manner, and multiple valid plans are maintained.

It must be ensured that the user is as free as possible while still experiencing dilemmas. In its current version the system is control-based. This means that the user selects actions until they choose to pass control back to the system, which then acts until a user action is required. When the user has control they can take any number of actions. The user can spend as long as they want considering their options.

When presented with a dilemma the potential consequences of each decision must be clear to the user before

they make their choice. Once they have chosen, these repercussions on the storyworld are implemented. The resultant state is thus entirely dependent on the user's decision.

The proportion of non-user dilemmas can be adjusted – by the story designer dependent on the genre or dynamically according to the frequency of user actions. The system is able to create a non-interactive story, meaning that there is always a story whether or not the user chooses to act. This adds to the illusion that these characters exist outside the user's scope. It also gives the user the option of not acting in the storyworld, whether for a long or brief period of time.

The system is able to provide direct responses to user actions through a system based on tit for tat reactions and utility scores. This involves a numerical utility value being assigned to each character in all story states. Actions change this value due to the corresponding change to the affected character's score. When the user acts in a way which affects the score of another character, that character responds by acting to change the user's score by the same amount. The use of utility values means that extension to additional actions requires only the association of a value with each. This method also makes system responses less predictable and more versatile. The responses update the state and thus effect the future path of the story - both immediately and in the longer term. These are an immediate effect of the user's actions and result in a story more specific to the particular user.

User Model Motivation

The user of an interactive drama system should be modelled rather than controlled. The story should adapt to the user's interactions rather than forcing the user to follow a particular storyline. Those choices which the user is likely to make should be identified by the model. By being combined with a fixed 'interestingness' value for each dilemma outcome this is currently used to select the next dilemma to be presented to the user. These values are based on the personal experiences and opinions of various soap viewers.

For example, there may be two dilemmas possible at a given stage. In one the user might have to decide whether or not to cheat on their partner - with an interestingness of 7 if they choose to do so, and of 4 if they choose not to. The other dilemma may require the user to choose between potential partners and have an interestingness value of 6 irrespective of the user choice. The user model then estimates the likely user choice and accordingly selects the next dilemma. So, for instance, if the user is expected to cheat on their partner then this dilemma will be the next to be presented to the user.

This could be extended to allow the system to search for the most interesting story path to a pre-defined fixed depth (dependent on the size of the search space and the speed of the search and planning algorithms). The user model would be employed to estimate the user choices on each dilemma. The expected total "interestingness" of that path could then be calculated. The system would then present that dilemma which has the highest chance of leading to the most dramatically interesting experience for the user.

This model could also be used to look ahead in planning. Rather than continually searching from the current stage, it can be determined which choices the user is likely to make and thus to plan from later stages. As a result responses to user actions would be more prompt. As the search becomes deeper this will become less accurate.

User Model Design

In order to achieve the aims of the user model the system must be able to accurately predict the user's decisions on presentation of a dilemma. This model is used to identify the decision the user is most likely to make when presented with a dilemma. The model is built up by observing previous user actions and making assumptions.

Each user is modelled according to various aspects and associated values. A series of rules are then used to predict future user dilemma decisions. The aspects modelled are: honesty, faithfulness, responsibility for actions, selfishness, preference for relationship or friendship, strength of character and morality. The value the user puts on their relationship and friendship with each character is also modelled, as is their strength of belief in each storyworld principle. Those aspects of the user to be modelled were selected for their generality and applicability to as many dilemmas as possible. These were chosen for their appropriateness to not only current system dilemmas but also to other soap dilemmas. In addition, important traits which are likely to affect decisions made in real soaps were considered. Although some of these may seem too similar, or perhaps too specific, they each apply to various different dilemmas, in a range of combinations and all are necessary.

Each modelled aspect has an associated integer value, which will change following observation of the user's behaviour. It was decided that each change to modelled values should be represented only as an increase or decrease of the associated value by 1. Since no single dilemma is more significant than any other – they are all required as components of the overall experience – this represents a more accurate model. Each of the aspects is initially assigned a value of 0.

Each dilemma updates certain criteria. These have been individually identified through observation of soaps and of users interacting with this system. An example can be seen in fig. 3. A rule associated with each dilemma reflects the balance of criteria values which will lead to each possible predicted user decision. An example of this is shown in fig. 4. If these values are equal, in the case of some dilemmas a second attempt at a prediction will be made (using less accurate criteria combinations) before a random prediction is made. It is clear that the performance of the model depends on the quality of these rules.

This model is fairly general. It is thus applicable to further dilemmas and will still be valid as the system develops.

Testing the User Model Accuracy

It will only be possible to utilise a user model if it is first determined that it is accurate. This means that testing is required. In this a user can be asked to interact with a system which does not apply the user model. The model is

```
Dilemma:  ''Would you like to cheat on
your partner X with Y?''

If yes:
decreased:
value for faithfulness
value for morality
value for relationship with X
increased:
value for relationship with Y

If no:
decreased:
value for relationship with Y
increased:
value for faithfulness
value for morality
value for relationship with X
```

Figure 3: This figure shows the updates to the user model which result from the user deciding whether or not to cheat on their partner. In this example the updates are symmetrical, although this is not always the case.

```
Dilemma:  ''Would you like to cheat on
your partner X with Y?''
negatives:
value for relationship with X
value for faithfulness
value for morality
positives:
value for relationship with Y
```

Figure 4: This figure shows the criteria used by the system in predicting the expected response of the user when asked to decide whether or not to cheat on their partner. This depends on the balance of values in the user model. If the sum of the positive criteria is greater than the sum of the negative criteria then the user will be expected by the model to choose 'yes', i.e. to cheat. If the values are equal a random prediction will be made.

still updated after each user decision. The model will be deemed accurate if it is able to predict the decisions the user makes when presented with a dilemma. It should continually evolve but will not be expected to be accurate for the first few dilemmas.

For each user dilemma, the system predicts their most likely decision. This means that the predictions should become more accurate as the story progresses. Each decision and prediction is written to file, along with the user model at the corresponding stage. This is then used to observe the accuracy of the predictions, and thus of the model.

Initially the model was tested by 3 users. Some of these tests involved longer stories – consisting of 8-10 dilemmas – for multiple narratives. The results indicated that the only predictions made incorrectly were in the early stages of the narrative when the model had insufficient information on the user.

It is intended to carry out more extensive evaluation of the user model with a larger number of users. In this care will have to be taken. Although modelling the user's strength of feeling for other characters is useful in predicting their dilemma decisions in the current system users often find it difficult, particularly initially, to believe in the characters as individuals. This could mean that their actions towards these characters are inconsistent with previous behaviour. In order to improve the user's belief in and familiarity with the characters, and thus to hopefully improve testing results, characters from existing soaps could be used. The user would then be able to select their soap and experience a story with characters from that soap. Of course these characters would not be modelled in any depth, but the model should be sufficient for the user to have a more enjoyable experience and for the model to be applicable.

Additional lack of familiarity with the system and their available options could prevent users in further testing from acting rationally and thus reduce the accuracy of the user model. This means that a large time commitment will be required by the users as they may need to build up familiarity before the model becomes fully accurate.

An evaluation of the effect application of the user model has on the user's experience of the story should also be carried out. The model will only be successful if it improves the user's experience in some way.

Extending the User Model

The credibility given to the user model could be adapted to depend on how recently the criterion being utilised was updated – since the user and their opinions are likely to change through the course of the interactive narrative.

Rather than having a definitive prediction the probability of the user making a particular dilemma decision could be approximated by the user model. The system could select that dilemma which had the highest chance of leading to the most dramatically interesting dilemma outcome.

An improved user model may result from the use of percentage value changes – rather than absolute – when a user model criteria is updated. This would take place within upper and lower limits and would reduce the bias effect a number of sequential similar dilemmas could cause to the model. As a result the user model could become more versatile and accurate.

A more accurate user model could result if user actions were modelled in addition to their dilemma decisions. This could be particularly the case in modelling their feelings towards other characters. For instance, if the user flirts with another character then it is likely that they have a higher value for their relationship with that character.

A model may be created which determines which dilemmas the user will find most interesting and thus presents these. The relevance of this is apparent when the use of the 'hook' in soaps is considered. If there is a storyline to which the viewer particularly relates then they are more likely to start watching that soap. Once they start, other storylines will begin to intrigue, if not to the same extent as that first

storyline. This then leads to regular viewing of that soap. However, more 'hooks' will be required to keep the viewer interested. This could be adapted for use in soaps through use of a user model which would determine storylines likely to be a 'hook' for the particular user (possibly through user categorisation). Regular 'hooks' (which should be more relevant over time) can then be incorporated.

In making actions available to the planner in the current system it is assumed that the user will act in a manner consistent with other soap characters and their individual traits, i.e. in accordance with the standard applicability check. An assumption model could be created to replace this applicability check with one more appropriate to the individual user, depending on their previous actions and decisions. This assumption model will develop and become more accurate as the story progresses. With use of this model the likelihood of the user acting consistently with the plan will be increased as the planner will have a more accurate idea of how they are likely to act. This in turn will increase the frequency and interestingness of dilemmas presented to the individual user.

Related Work

Only two other interactive drama systems utilise a user model. The IDA (Magerko 2005) user model is used only to direct the user within the story's pre-defined overall plot structure. In IDtension (Szilas 2003) the user takes turns with the system to choose actions for the story as a whole. If they are modelled to consistently choose actions which avoid violence, the system can present them with a dilemma in which they must choose a violent action in order to achieve the pre-defined goals of the story. The dilemmas here are for the user as an external observer of the system, rather than as a character.

Conclusions

In this paper we presented an interactive narrative generator that is able to create long, and potentially infinite, story lines that incorporate dilemmas to add dramatic tension. The stories are dynamically created based on user decisions and actions as well as adapting to the user's tendencies. The user model presented is able to predict user decisions on dilemmas, and evolves throughout the experience. Being able to predict likely user responses to dilemmas means that the system has a fairly good idea of what the user is like and can thus act in a way to optimise the experience of that user. The system can also present dilemmas which are likely to have the most interesting consequences, in accordance with expected user decisions.

The use of all soap dilemmas to determine appropriate modelling criteria – as opposed to only those dilemmas currently available to the system – should mean that the model remains relevant when the number of dilemmas is increased, and that no additional criteria will be necessary. When applying to additional domains further investigation of criteria will be required.

The user's experience and opinions thereof when using the system both with and without the user model should be analysed and compared in order to determine whether the user model improves the user's experience, as is intended.

Acknowledgements

Particular thanks are due to George Barber for his knowledge of soaps and ability to identify dilemmas therein; and to James Barber and Francis Barber for their hours of testing.

References

Bates, J. 1992. Virtual reality, art, and entertainment. *Presence: The Journal of Teleoperators and Virtual Environments* 1.

Cavazza, M., and Charles, F. 2002. Character-based interactive storytelling. *IEEE Intelligent Systems* 17.

Crawford, C. 2004. *Chris Crawford on Interactive Storytelling*. New Riders.

Fairclough, C. 2004. *Story Games and the OPIATE System*. Ph.D. Dissertation, University of Dublin - Trinity College.

Field, S. 1984. *The Screen-writer's Workbook*. New York: Dell Publishing.

Karlsson, B.; Ciarlini, A. E. M.; Feijó, B.; and Furtado, A. L. 2006. Applying a plan-recognition/plan-generation paradigm to interactive storytelling. In *Workshop on AI Planning for Computer Games and Synthetic Characters*.

Magerko, B. 2005. Story representation and interactive drama. In *1st Artificial Intelligence and Interactive Digital Entertainment Conference*.

Mateas, M., and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. *Game Developers Conference, Game Design track*.

Rousseau, D., and Hayes-Roth, B. 1998. A social-psychological model for synthetic actors. In *International Conference on Autonomous Agents*.

Sgouros, N. M. 1997. Dynamic, user-centered resolution in interactive stories. *International Joint Conference on Artificial Intelligence* 2.

Szilas, N. 2003. Idtension: a narrative engine for interactive drama. In *1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment*.

Thomas, J. M., and Young, R. M. 2006. Author in the loop: Using mixed-initiative planning to improve interactive narrative. In *Workshop on AI Planning for Computer Games and Synthetic Characters*.

Young, R. M. 2001. An overview of the mimesis architecture: Integrating intelligent narrative control into an existing gaming environment. In *The Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*.