Article:
Barber, H. and Kudenko, D. orcid.org/0000-0003-3359-3255 (2009) Generation of Adaptive Dilemma-Based Interactive Narratives. Computational Intelligence and AI in Games, IEEE Transactions on. 309 -326. ISSN 1943-068X

https://doi.org/10.1109/TCIAIG.2009.2037925

Generation of Adaptive Dilemma-based Interactive Narratives

Heather Barber and Daniel Kudenko

***Abstract*—The Generator of Adaptive Dilemma-based Interactive Narratives (GADIN) presented in this paper dynamically generates interactive narratives which are focused on dilemmas to create dramatic tension. The system is provided with knowledge of generic story actions and dilemmas based on those clichés encountered in many storytelling domains. The domain designer is only required to provide domain specific information, for example regarding characters and their relationships, locations and actions. A planner creates sequences of actions that each lead to a dilemma for a character (who can be the user). The user interacts with the storyworld by making decisions on relevant dilemmas and by freely choosing their own actions. Using this input the system chooses and adapts future story lines according to the user's past behaviour. Previous interactive narrative systems often have content creation and ordering requirements which restrict the possibility for sustaining the dramatic interest of the narrative over a long time period. In addition, many of these systems are not easily transferable between domains. In this paper the GADIN system is demonstrated to both be able to maintain the dramatic interest of generated narratives over a long time period and to have a core architecture which is applicable to any domain.**

I. Introduction

STORYTELLING is appreciated by many people, as both teller and audience. In the past stories were only told orally, with audience participation. It is still true that listening to a friend narrating a story is an enjoyable way to spend time. However, as storytelling has evolved - through drama, writing, print, film and television - interactivity has been neglected. Receivers (listeners, readers or viewers) of a story will often want to become more involved in the storyworld, perhaps even to become a character. An interactive narrative offers a world in which the participant can have a real effect - both long and short term - on the narrative which they are experiencing.

Most modern computer games involve a story, which in most cases is an essentially linear story or series of stories. As a result this element usually violates a basic requirement for such games - the need for interaction which has a clear effect. There are games with no explicit story structure - in which the player is encouraged to perceive their own stories within the world. These stories are truly interactive but lack the skill of a playwright and subsequent high level of dramatic interest. An interactive narrative combines the free interactions of the player with this play-writing skill to create a dramatically interesting game playing experience.

There are various definitions and conceptions of interactive narrative [1]–[11]. These have core similarities and identify the same essential requirements. Having considered these definitions, interactive narrative as it will be considered in this research can be defined: *An interactive narrative is a game world in which the user-controlled character(s) can physically and socially interact with ideally (perceived) total freedom while experiencing a dramatically interesting narrative which is fundamentally different on nearly every play - dependent on the user's actions.*

There are many existing interactive narrative systems [1]–[8], [12]–[16]. While these systems advanced the area in several aspects, many challenges still remain. The GADIN (Generator of Adaptive Dilemma-based Interactive Narratives) system addresses (amongst others) two of the open challenges: maintaining the dramatic interest of the narrative over a longer period, and domain independence – the ease of transfer to new story domains.

Many storytelling genres make frequent use of clichéd storylines which are created around dilemmas to storyworld characters. These dilemmas can be generalised and the GADIN architecture uses planning to achieve dilemmas, the combination of plan and dilemma constituting a dramatically interesting substory of the generated narrative. Characters act and experience these dilemmas in the course of the narrative, making decisions and action choices depending on their individual properties and state. The GADIN architecture is discussed further in section II. In the interactive version (section II-G) the user is able to freely select their own actions and will also experience these dilemmas, which can be dynamically selected to involve the most difficult decision for the individual user. An example GADIN experience is given in appendix A.

The application of GADIN to specific domains is discussed in section III. The current applications are to a soap world (section III-A) and to a children's dinosaur adventure (section III-B). An interactive narrative experience such as GADIN is a game in its own right, and can be played as a stand-alone experience. However, the GADIN technology can also be utilised as a component within more traditional computer game genres. For example in an exploratory game, such as *Fable* [17], a player could enter a town and could experience an original narrative within that town. This narrative would probably not have an effect on the overall game as there is no prescribed way to experience a GADIN interactive narrative. In a game such as *Grand Theft Auto* [18] it may be possible for the player to experience a GADIN narrative as a mission. The advantage would be the experience of a series of completely original narratives. This would potentially be extendible to any games which use missions, or quests, and to all such missions. Within the structure of the GADIN system such quests may be relationship-based, or more commonly mission-based. This will result in missions which are truly unique

Department of Computer Science, University of York, Heslington, York, YO10 5DD.

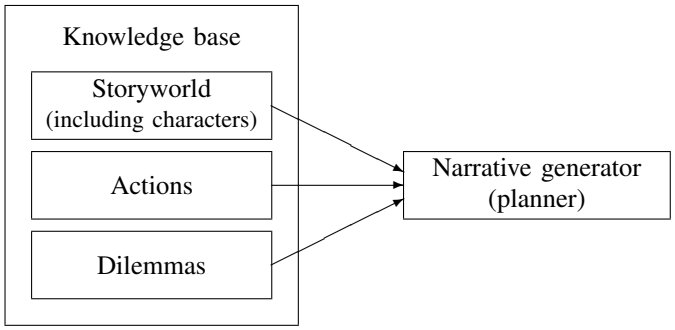

Fig. 1. An overview of the GADIN architecture

for each individual player, and which have an original and interesting outcome.

In section IV the methodologies used to evaluate the GADIN system are detailed. The criteria of scalability of an interactive narrative and ease of transferability to a new domain are discussed in more detail in section V. These are discussed in the context of the main existing interactive narrative systems. The GADIN system is shown to achieve both scalability and transferability. The outlook for the system is considered in section VI.

II. The GADIN architecture

Figure 1 shows an overview of the GADIN architecture. The GADIN knowledge base consists of: the storyworld (which includes information regarding the characters); story actions in which the characters can participate; 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 utilised in the generation of a narrative. In the GADIN system planning is used to achieve the preconditions of dilemmas through story actions, dependent on information in the storyworld database.

A. Storyworld

The storyworld is the world in which the narrative will be generated. It consists of all of the components which must be utilised in the creation of a narrative. There must be: characters who can act within the world; at least one location at which the narrative can take place; and objects which can be interacted with. The aspects of characters, locations and objects available within a GADIN storyworld are discussed here.

Characters: Each character's potentially associated traits include: attributes, characteristics, personalities, principles, aspirations, skills and dispositions. A character's associated attributes can include information such as attractiveness and gender. These are generally physical traits of a character. Characteristics are slightly more variable than attributes, and are usually features relating to the mentality of a character, for example generosity and morality. A range of values is associated with each attribute and characteristic.

Where required it is possible to specify genre specific character personalities not fully deducible from other character traits, such as being a fundamentally evil character. Characters hold storyworld principles, such as monogamy, to make their behaviour more believable. Under specified pressures and circumstances, principles may have their associated strength of belief changed. Characters have aspirations, for example wanting a baby, and skills, such as being able to swim.

Characters have an associated disposition, which is defined along each of a number of dimensions. Before the narrative generation begins each character is randomly assigned a value for each disposition dimension. Throughout the narrative characters will choose to take actions which are consistent with their disposition. For example a charming character may consistently tell others they look nice.

The specifics of character traits can be determined by the storyworld designer. The traits and associated values in the soap version of the GADIN system are shown in table I.

Characters have storyworld relationships with one another, for example friendship, love and familial relationships. Relationships are unidirectional and have an associated strength, although feelings of one character for another affect the reciprocity. In the current system relationships only exist or do not.

Locations: A series of genre-specific locations can be defined for each storyworld. At any given time in the narrative it is possible for each character to be at only one of these locations. Direct interactions between characters will only take place if they are at the same location or able to communicate in some way, such as on the telephone (if this would be appropriate for the domain).

Each location has a name, and a series of values which determine its nature. These may include whether or not the location is: appropriate for a party; a place where it is possible to steal; a place where it is possible to drink; or a place where drinks can be bought (for others). Skills required of characters, such as being able to swim, before they can move to a location can also be defined.

Objects: Where appropriate for the domain there will also be objects within the storyworld. Objects can be attached to an owner, after which they will belong to that owner and move with them. Objects which can be stolen will have a 'true owner'. If a character steals such an object then they become the owner, and can treat the object as such, but the object may be required to be returned to the true owner. If the object is given away or sold then the true ownership will change.

B. Actions

Aristotle [19] observes that "life consists in action" and therefore within a drama "character comes in as subsidiary to the actions". To generate a narrative those actions which can take place within the storyworld must first be defined. The requirements for GADIN actions are discussed in this section.

STRIPS formulation: The actions which can occur in the course of a narrative must be specified for each domain. There may also be actions which will be appropriate in many domains. Every possible action should be included for each genre. Each action will have associated: conditions which must be satisfied before execution (preconditions); and effects representing changes to the storyworld following successful

TABLE I
THIS TABLE SHOWS THE TRAITS AND ASSOCIATED VALUES USED IN THE SOAP VERSION OF GADIN.

Category	Trait	Values
Attribute	Attractiveness	-1, 0, 1
Attribute	Gender	male, female
Attribute	Age	child, teen, young, middle aged, old
Attribute	Sexuality	homosexual, heterosexual, bisexual
Characteristic	Generosity	-1, 0, 1
Characteristic	Morality	-1, 0, 1
Characteristic	Selfishness	-1, 0, 1
Personality	Personality type	bad_boy, busy_body
Principle	Not stealing	true, false
Principle	Not drugging others	true, false
Aspiration	Wanting a baby	true, false

Operation: `Move between location l and location k`
Preconditions: `at(l) ∧ path(l , k)`
Effects: `at(k) ∧ ¬at(l)`

Fig. 2. An example of a STRIPS-representation action.

execution. Preconditions and effects can include any proposition which has been defined within the domain. Effects can also be changes to character dispositions.

For example the action of a character moving between locations `l` and `k` has preconditions of the character being at location `l` and there existing a path between locations `l` and `k`. The effects of this action are that the character is at location `k` and is no longer at location `l`. This follows the STRIPS representation and is shown in this form in figure 2.

Applicability check: Before an action is made available to the system for use within a storyline an applicability check is performed on the involved character's traits and disposition. An action can only be utilised if its applicability is high enough for the acting character – to ensure that the action is of the type that the acting character would make. This check is supplementary to the preconditions of an action and incorporates conditions which cannot be specified through use of STRIPS-style preconditions. This helps to ensure that each character acts in a manner which is consistent with their traits and how they have acted previously, while at the same time avoiding predictability.

An example in the soap version of GADIN (and therefore based on actions which have been observed to occur in soaps) is a character X starting to fancy another, Y. The precondition is that X does not fancy Y and the effect is that X does fancy Y. This action would not be appropriate for some characters, for instance a particularly attractive character would not start to fancy a very unattractive character (unless there are other incentives, e.g. of a financial nature). The applicability check ensures that each action is appropriate for the acting character, in this case it should be ensured that Y has attractiveness $\geq$ X's attractiveness. There may be other applicability conditions utilised for this same action.

Adverbs: Adverbs can be associated with character actions. The adverb selected as an action descriptor will be randomly chosen from those associated with the disposition dimension which has the greatest absolute value for that character in the current storyworld state. Modifying adverbs can be used if these dispositions have very high or low values associated. Character dispositions will become clear from the actions they choose and the manner in which they carry out these actions. For example a character action may be that they are going to `move to the forest`. If their disposition has the greatest value associated with the happiness dimension, an adverb randomly selected from the appropriate possibilities could be `joyously`, and the character would thus `move to the forest joyously`.

Utilities: Each character has an associated utility (or score) in each storyworld state. This is changed as the state of the storyworld changes through character actions. This reflects the assumed positives and negatives of that state for each character. It is calculated appropriately for the domain, for example in the soap domain a character will have a higher utility if another character fancies them and characters who want a baby will have a higher utility in states in which they are pregnant. The utility of a character in a storyworld state is the sum of the utilities associated with everything that holds for that character in that state.

C. Dilemmas

Narratives often centralise on clichéd plotlines. These will generally culminate in a dilemma involving the main character. This is not true in all storytelling domains, but is found to be the case in a wide range of domains, for example 'chick flicks', James Bond-style adventure films and soap operas (soaps). Within these genres writers will utilise such clichéd plotlines in the creation of narratives, building up the narrative around them in the form of conflicts, or dilemmas, to characters.

A range of such dilemmas can be identified and generalised. Once the general form of each dilemma has been determined, it is possible for a computerised storywriter to generate a narrative around these. Here the narrative is built around the cliché, and it is the cliché as well as the narrative which the audience appreciate, the very repetitiveness and familiarity of the dilemmas adding to the dramatic interest. The generation of a narrative in this manner provides a story which is original each time, as the variations in the clichés will result in new narrative.

A decision on a dilemma involves only two recipients of direct differing utility payoffs. Other dilemmas can be reduced to this form. Five such dilemma categories were identified.

These consist of all payoff matrices with two recipients where there is a dilemma involved. This may require characters to be friends or enemies and where relevant this is stated with the dilemma utility matrices. The relevant categories are: Betrayal, Sacrifice, Greater Good, Take Down and Favour.

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 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 a Betrayal dilemma.

Sacrifice: A character involved in a Sacrifice dilemma is able to choose an action which will result in their lowest possible utility but also the best outcome for their friend. These characters must be friends for this to be a dilemma.

An example of a Sacrifice dilemma occurs when a character has committed a crime which their friend has been accused of. 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 highest 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 a Greater Good dilemma involves a character deciding whether to give something (such as information or a friend) to their enemy to save themselves.

Take Down: In a Take Down dilemma a character has the option of an action which will result in their lowest possible utility but also the worst outcome for their enemy. The characters must be enemies for the dilemma to exist.

A character deciding whether to harm (or even kill) their enemy in full awareness that they will be punished for this is involved in a Take Down dilemma.

Favour: A favour dilemma causes a character X to have to choose between two actions where there will not be any immediate discernible benefit to X as a result of their decision. The utilities of characters Y and Z will change as a result of this action choice. 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. When presented with a favour dilemma the character making the decision will not receive any direct utility from their action regardless of their choice, nor will there be any discernible benefit to the character making the decision of choosing one character over the other.

An instance of a Favour dilemma occurs when a character must choose between potential partners.

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 belongs to 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 domains. From these categories dilemma instances can be found and generalised within each domain. From the generalised form of the dilemma, the system will be able to create new dilemmas. In the presentation of these original narratives are generated.

D. Narrative generator

Dilemmas require characters to make fundamentally difficult decisions within the course of the narrative and thus create dramatic interest. 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 planner to achieve these preconditions, using actions which are possible within the storyworld, and thus to enable presentation of dilemmas. Such a plan constitutes the build-up - the essence of the story itself - and becomes a storyline when presented. The narrative will be made up of a series of such sub-stories, dynamically selected according to dramatic interest.

On being passed a dilemma the planner (which is based on the GraphPlan algorithm [20]) finds all plans to achieve this dilemma given the current storyworld state and background knowledge. A planning thread constantly finds possible plans for all available dilemmas. Once the system is ready to present a dilemma it selects one of the currently available plans. The sequence in which the dilemmas are selected for presentation must depend on what has happened previously to become part of a consistent story. Following presentation of a dilemma the next most appropriate must be selected and its presentation will be attempted. The most appropriate dilemma is selected depending on the previous dilemmas which have been presented and the frequency of dilemma use. The planner cycles constantly through all dilemmas, searching from the updated state to find any possible plans for dilemmas. The plan is presented as a sequence of actions prior to a dilemma - for which the decision and outcome are shown.

It is important that good storylines (and corresponding dilemmas) do not become devalued by overuse. Each dilemma thus has an associated frequency rating which reflects its frequency in other narratives in the current domain. If the number of times a dilemma has occurred in the recent narrative history is greater than the dilemma frequency rating then this dilemma cannot be presented. An additional constraint ensures that exactly the same dilemma will not be experienced by the same character more than once within a certain period of time.

The potential consequences of each decision must be clear to the deciding character before they make their choice. Once a choice has been made, the system will update the storyworld state in accordance with that choice. The system then plans from the new state to be able to present another dilemma - thus continuing the interactive narrative. This sequence of events is demonstrated in figure 3.

When making decisions on dilemmas characters will act in accordance with their individual traits and circumstances. For example if they are married then they are less likely to have an affair.

E. Responding to dilemmas

Following presentation of a dilemma there will be immediate changes to the storyworld, the effects of the dilemma

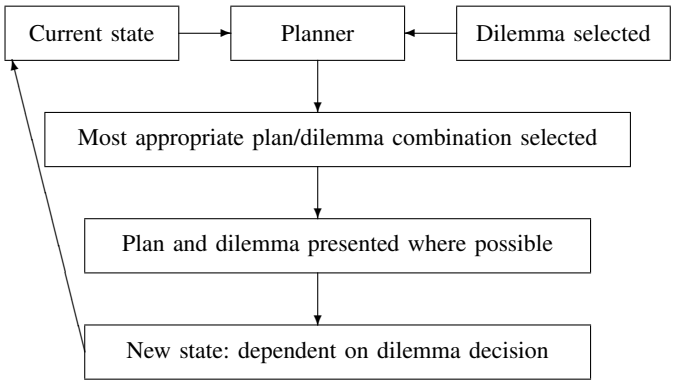

Fig. 3. An overview of the system moving between states dependent on plans, dilemmas and decisions.

decision. In addition there will be responses from other involved characters to the deciding character. The specifics of the response depend on the individuals involved and their traits, dispositions and relationships as well as the category of the dilemma.

To provide these responses dilemma decisions change the utility values of affected characters in an author-defined (and potentially character dependent) manner. These characters will thus act to change the deciding character's utility in the appropriate corresponding manner. Each action changes the state and thus the utilities of characters.

F. Example dilemma and plan

In this section an example of a dilemma is discussed. Basic information for an imaginary storyworld state is given and the creation of a plan is described.

The plotline of a character being presented with a dilemma involving cheating on their partner is frequently used in soaps. This dilemma can be categorised as being of type Betrayal. Its general form is:

$\mathbf{A_X}$**:** `cheat on partner (X)`
preconditions: `partners(X,Y)` $\wedge$ `fancies(X,Z)` $\wedge$ `fancies(Z,X)`
dilemma (to character X): ``Would you like to cheat on your partner character Y with character Z who fancies you?''
if X chooses to cheat:
`add to state: cheating(X,Y,Z)`
if X chooses not to cheat:
`delete from state: fancies(X,Z)`

A possible soap action involves X starting to fancy Y, which has preconditions:

`fancies(Y,X)`
`not(fancies(X,Y))`

and effect:

`fancies(X,Y)`

In STRIPS form this is:
Operation: `starts_to_fancy(X,Y)`
Preconditions: `fancies(Y,X)` $\wedge$ $\neg$`fancies(X,Y)`

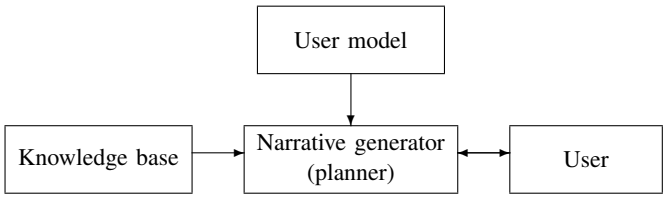

Fig. 4. An overview of the interactive GADIN architecture.

Effects: `fancies(X,Y)`
One of the circumstances in which this action is considered applicable is when `attractiveness(Y)` is greater than `attractiveness(X)`.

In this imaginary storyworld there are three characters: Jim, Tina (both with attractiveness 0) and Liz (who has attractiveness 1). In the current state Jim and Tina are partners (`partners(jim,tina)`) and Liz fancies Jim (`fancies(liz,jim)`) but this is not reciprocated ($\neg$`fancies(jim,liz)`).

The planner takes the current state and applicable actions and attempts to achieve all of the preconditions of the dilemma. With X = `jim`, Y = `tina`, Z = `liz` two of the preconditions of the dilemma are satisfied but the third is not. Since the action involving Jim starting to fancy Liz is applicable in this case the planner can use this action and thus achieve all of the preconditions of the dilemma. The one act plan is thus: `starts_to_fancy(jim,liz)`. If this dilemma is found to be appropriate then its plan will be presented. The effect will be that the new state will contain the elements: `partners(jim,tina)` $\wedge$ `fancies(jim,liz)` $\wedge$ `fancies(liz,jim)`, and this instantiated Betrayal dilemma can thus be presented to Jim.

G. Interactivity

It is possible for the user to become a character in a GADIN experience. They are able to act freely and a dramatically interesting narrative will still result. An overview of the interactive GADIN architecture can be seen in figure 4. In this a user model may also be included (as shown). This is an optional component which can be employed to ensure that the narrative's dramatic interest is maximised for the user. The user model is discussed further in section II-J.

User dilemmas: In the interactive GADIN experience the user becomes a character within the storyworld. Within the course of their experience the user encounters dilemmas which require them to make fundamentally difficult decisions. 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 entirely dependent on the user's decision. It is ensured that the user experiences a reasonable proportion and balance of dilemmas while the overall frequency is as would be expected for the genre.

User actions: Every action that other characters within the system can take is available to the user who is able to freely specify their own actions within the scope of the current genre.

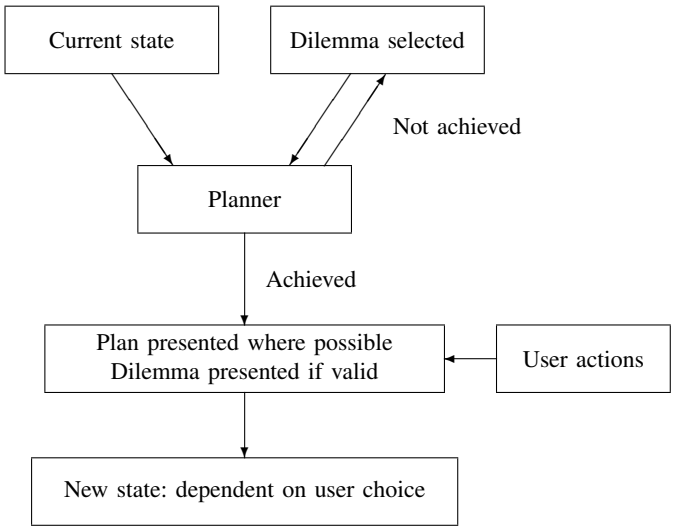

Fig. 5. An overview of the interactive version of the system moving between states dependent on plans, dilemmas and user decisions.

An overview of the narrative generator which incorporates user actions is shown in figure 5.

Once a plan has been chosen it is then presented to the user in such a way as to incorporate their actions. Each plan action can be alternated with a user action, or all plan actions at a particular level can be presented to the user before allowing them to act – depending on which method is the most appropriate for the domain. If the user acts in a manner which satisfies the necessary preconditions of actions at the next required level of the plan then the presentation of the plan will continue. As soon as it becomes possible to present the dilemma this is done.

In the planner it is assumed that the user will act consistently with the manner in which characters with similar traits would act within the current domain. In presenting the plan the user may not take these actions. The user is entirely free to select their own actions and will not know the plan which the system is attempting to follow. If the user acts in a way which violates the plan the system will be required to select another plan. Where possible plans should be created with which the user is more likely to act in accordance.

It must be ensured that the user is as free as possible while still experiencing dilemmas. Any user action which satisfies the preconditions of the next stage of the plan is acceptable, but even then the user has a wide range of options and may not act as required by the plan. To help overcome this problem shorter plans are favoured. This means that there are less opportunities for the user to act outside the plan, while still creating plans in which their actions will have an effect. Narratives of the same length will involve more drama if plotlines are shorter.

An attempt is made to coerce the user into acting in the way required by the current plan. For example if it is required that the user moves from location `l` to location `k` their friend can go to location `l` and ask the user to join them in going to location `k`. A record of coercions is maintained and used to ensure that the user will never be repeatedly coerced for the same action or group of actions.

In its current version GADIN is control-based. This means that the user selects actions until they choose to pass control back to the system, which – depending on the domain – either: acts until a user action is required to satisfy required preconditions; or allows a single character action or dilemma. When the user has control they can take any number of actions. The user can spend as long as they choose to considering their options.

It is the user's decision how to respond to dilemmas. They can ignore dilemma decisions or react to them in any manner which they consider appropriate. Other characters will react to the user's dilemma decisions in the manner discussed in section II-E. When the plan for a dilemma to another character requires user actions these can be attempted to be incorporated in the same way as for plans for user dilemmas.

Affecting characters: The user is able to act in a way which could cause another character to experience a dilemma. If they do so then that character will immediately be presented with the appropriate dilemma. The user may cause changes to the disposition of another character, and thus how that character will act in the following stages of the narrative. For example in the dinosaur adventure version of GADIN if the user plays with another character the value associated with the outgoingness dimension of their disposition is increased.

When the user selects an action which does not correspond to the current plan a utility-based response to their actions is given. This method is used as identifying patterns in large numbers of user actions is complex and requiring this would reduce the extendibility of the system. When the user acts in a way which affects the utility of another character (due to the resulting changes in the storyworld state), that character responds by acting to change the user's utility by the same amount. The system computes user utilities in the same way as the utilities of other characters. If the user's actions have not changed the utility values of any other characters then there is either no response or a response which is deemed to be the most appropriate, dependent on the user's actions and how they have affected the other characters. An example would occur when the user stops fancying a character and thus reduces their utility. The response would be randomly selected depending on the traits and circumstances of the responding character. It could be that the character responds by ceasing fancying of the user (if this is possible), or that the character feels rejected and thus encourages the user to betray their principle and to steal. The use of utility values means that extension to additional actions and new domains requires only the association of a value with each. This method makes system responses less predictable and more versatile.

H. Events

An event can be initiated by one or more characters and involves at least those characters. Any characters, including the user, may initiate or be involved in events providing the preconditions have been satisfied and there is sufficient motivation. Events are dissimilar to actions in that they directly involve more than one character. Events can cause the

occurrence of dilemmas to be more likely or result from the outcome of dilemmas. The effects of a dilemma will depend on their outcome.

Examples of possible events include: proposals, weddings (including their interruption), funerals and parties. The motivation for a proposal could be that a character wishes to compensate for having had an affair. An example of an event causing a dilemma would be a party which causes two characters to be at the same location where they have the opportunity to start an affair.

I. Knowledge

In conventional storytelling information is often revealed to the audience but not to the characters to give the audience a sense of suspense. The audience will know more than the characters and can use their knowledge to further interpret and understand the narrative. If the user of an interactive narrative is told as much as the audience yet is simultaneously a character the manner in which they act will be affected and the dramatic interest may be reduced. For example if there is a murder committed and the user, as audience, has knowledge of this murder then they will act differently as a character.

It was decided to enable certain actions and information to be concealed, such as the identity of a murderer. This means that it is necessary to record some of the information characters, and the user, have about the storyworld. This is referred to as their 'knowledge'. Each character (including the user) has a hidden list of knowledge associated. This changes when they make up or are told a piece of knowledge, or when it is revealed that the knowledge was not true.

For example a character may see another near the scene of a murder and assume that they were the murderer. They will then add this knowledge to their list of knowledge. If a character has an item in their list of knowledge they can tell it to another character, and it will subsequently be in that other character's knowledge list. If an item is publically revealed to be untrue it is removed from all characters' knowledge lists. Any item which characters are told they assume to be true and add to their knowledge list. Characters having items in their knowledge list can be a precondition or effect of both actions and dilemmas where required.

J. User model

The GADIN system is able to create a model of the user based on their dilemma decisions and action choices and use this to select future dilemmas to be presented to them – depending on which are likely to be the most conflicting. To achieve this the system must accurately predict the user's decisions on presentation of a dilemma.

Each user is modelled according to various aspects and associated values. The specific aspects of the user which are modelled are determined depending on the domain. Examples include honesty, responsibility for actions, and strength of character. The value the user puts on their relationships with other characters and storyworld principles are also modelled. Those aspects of the user to be modelled are selected for their generality and applicability to as many dilemmas as possible.

```
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
```

Fig. 6. 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.

Each modelled aspect has an associated integer value, which changes following observation of the user's behaviour.

Updates to the values associated with each of the user model criteria are made following each user action and dilemma decision. Each dilemma updates certain criteria, which should be specified in defining each. Since no single dilemma is more significant than any other – they are all required as components of the overall experience – each relevant aspect is updated by the same percentage regardless of the dilemma and its category. The percentage update is inversely proportional to the difference in the number of positive and negative updates to that criteria. Following each user action the values associated with relevant criteria in the user model are updated. For example if the user flirts with another character then it is likely that they have a higher value for their relationship with that character.

An example of the user model updates in the soap domain following a dilemma decision can be seen in figure 6.

Those choices which the user is likely to make are identified by the model. A rule associated with each dilemma reflects the balance of criteria values which will lead to each possible predicted user decision. The performance of the model depends on the quality of these rules. Each dilemma will appeal to a different type of user - the user model chooses those that are likely to appeal to the current user. The dilemmas which are most difficult to predict are likely to involve the most difficult decisions for the user and thus greater conflict in the narrative. This can be determined by considering the balance of criteria values. That dilemma which has the smallest difference in the values associated is the most difficult to predict.

An example of this might occur in a dilemma in which the user has a friend who is encouraging them to betray a principle. The balance will involve the strength of friendship against the strength of belief in the principle. If the user holds this principle particularly strongly but also truly values that friendship the decision will be particularly difficult for them to make. An accurate user model will be able to predict this

```
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
```

Fig. 7. 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. The difference between the sum of the positive and the sum of the negative criteria will determine how difficult this decision is expected to be for a particular user at a given stage in the narrative.

and thus present such a dilemma to the user.

An example of the prediction balance for a dilemma in the soap version of GADIN is shown in figure 7.

K. Overcoming planning problems

In some storytelling domains it is necessary for the narrative to continue indefinitely, or for a long time period – for example in soaps. It is essential that the dramatic interest is maintained throughout the narrative. A number of techniques can be used to ensure that the narrative continues for as long as is required for the domain, and infinitely if necessary. These are discussed in this section.

If no valid plans have been found, but at least 3 attempts have been made, some of the applicability restrictions may be relaxed. The domain creator will determine the extent of this, and should ensure that the relaxation is never too extreme. An example might be that if characters are normally unable to start liking each other in plans they could be able to in relaxed applicability plans. This allows for the incorporation of those actions which are consistent with the narrative only if they occur less frequently than standard actions. It is acceptable for characters to occasionally act less strictly in accordance with their individual traits, and this can make the narrative more interesting. The relaxation of applicability may be layered, with each layer involving a greater degree of relaxation, depending on the domain. This is implemented through use of a variable, which determines whether relaxed applicability is used in the planning thread, and which will be changed as soon as presentation of a plan commences.

When no action has taken place for a long period of time (by default 60 seconds, normally the response time is just a few seconds) it is possible for characters to act randomly towards each other or the user. This enables the narrative to continue without waiting for time to be spent searching for a plan which may or may not contain this random action. Once one such action has taken place the next to be selected will, if possible, be a utility-based response to the action (and the character) taking the previous action (using the method discussed in section II-E). This will only happen once, to reduce the likelihood of the narrative deteriorating into an interplay between two characters – unless one of the involved characters is the user, when there will be a response whenever possible. The non-deterministic character actions never involve a major utility change, as there is no motive for this. The potential for a deciding character to change their own utility is allowed for, although no response is made to this.

L. Finite narrative

Narratives with a finite plot structure conventionally have an ending which is clear, satisfactory and understandable. For generation of a finite narrative in GADIN a storygoal is selected randomly from everything which is not true in the initial state of the storyworld, but which could be true within the current domain. For example the user having possession of a certain object. The satisfaction of a storygoal signifies the end of the narrative. The outline of the finite narrative generation process is shown in figure 8.

Given actions (including those for the user) and dilemmas within the storyworld GADIN creates a storyplan which satisfies the storygoal. The storyplan will be followed appropriately throughout the narrative. If the success of the current storyplan becomes unlikely GADIN can replan for the same storygoal (with a different plan) or identify and plan for a new storygoal. The ability to dynamically select a new storygoal gives the user a clear effect of their actions on the long-term of the narrative. Once a new storygoal has been selected the previous storygoals will still be maintained as possible endings for the current narrative.

For the ending of the narrative to be clear to the user they must know the storygoal. A character at the same location as the user (one will move there if necessary) will hint at the storygoal to the user, for example telling them that `Going to the cave is good` if the storygoal requires that the user be at the cave.

M. Narrative generation process

In figure 9 an overview of the main narrative generation process is given. Inclusion of each of the stages depends on the domain and the requirements of the creator. For example if only the user experiences dilemmas in the generated narrative then the stages which check for dilemmas for other characters (ii and A) are excluded. This process will continue until the storygoal has been achieved (in a finite narrative) or the user chooses to end their narrative experience (in an infinite narrative).

III. Application domains

This section discusses the application of the GADIN architecture to specific narrative domains. In section III-A the application of GADIN to a soap world is discussed, and a GADIN dinosaur adventure is detailed in section III-B.

A. An interactive soap

This section discusses the application of the GADIN system to a soap domain. Soap operas (or soaps) are a popular television and radio means of entertainment. There is frequent use of clichéd storylines in soaps. The infinite nature of soaps means that there is no overall plot structure but rather an infinite series

```
The user is able to act first. Whenever the user returns control to the system it attempts
the listed possibilities until one succeeds, after which control will be returned to the user.
After each user or system action the state is checked to see if any existing storygoal has been
satisfied, with this process continuing if not.
  1) If a dilemma has previously been presented within this experience present the next action
     of the storyplan.
  2) If the success of the storyplan has become unlikely find a new storyplan (for an existing
     storygoal or (if this is not possible) a new storygoal).
  3) Check if there are any dilemmas available for presentation to the user, and if so present
     as much of the plan as possible, or the dilemma.
  4) Check if there are any dilemmas available for presentation to other characters, and if so
     present as much of the plan as possible, or the dilemma.
  5) Allow a character to select an appropriate action depending on their individualities and
     relationships.
```

Fig. 8. An outline of the narrative generation process for a finite narrative in GADIN.

of 'mini-stories'. Soaps involve characters similar to the target audience experiencing their (very dramatic) everyday lives. There are many English soaps but those which are particularly focused on in the soap version of GADIN are *Hollyoaks*, *Eastenders*, *Coronation Street*, *Emmerdale* and *The Archers*. In identifying dilemmas from soaps it was found that they fell into only three of the five possible categories, namely: Betrayal, Sacrifice and Favour.

An extract from one user's experience with the system is given in Appendix A. Lines 2-26 show the user interacting with the plan for a dilemma, and subsequently being presented with that dilemma. A number of other dilemmas are presented to the user throughout this extract. One of the dilemmas to characters other than the user is on lines 93-101, with a plan not involving any user actions. The user acts in a manner which causes a character dilemma on lines 71-77. On lines 37-41 the system responds to the user in a manner unrelated to a specific plan. If the user does not cooperate with the plan, as on lines 82-85, the system continues the narrative with the next dilemma and corresponding plan.

B. A dinosaur adventure

This section discusses the application of the GADIN system to a finite children's short story. In this the user interacts in a dinosaur world in which the other characters are all dinosaurs. The user is able to interact with the dinosaurs – and objects in the world – in the ways which they might expect to in a children's story.

In the dinosaur adventure the user begins by being transported to the prehistoric world. Although this is an inevitable beginning, it is simply a device which takes the user into the storyworld. When the user achieves a storygoal they are taken back to their original world. There is also predictability in this return to the original world, but the manner in which the user returns is changeable (depending on the storygoal) which adds variety and interest to the ending of the narrative. The sense of inevitability, that they will always get home at the end, is common in children's literature and will be necessary to provide them with a happy ending. This is in accordance with Aaron Shepard's requirement that in a children's story there must be "a problem the main character must resolve" [21] and is not an essential requirement of finite narrative generation in the GADIN system.

The dinosaurs (storyworld characters) have dispositions which take a value along each of 3 dimensions. The dimensions are happiness, outgoingness and agility. As this is a children's story the narrative focuses on the main character and only the user experiences dilemmas.

Assumptions are made as to whether the user likes or does not like other characters. This represents the user's feelings without requiring them to be explicitly expressed. Dinosaurs never express their dislike of each other or the user. Instead characters will `chase` others to express their dislike. This is represented internally by the GADIN system as no longer liking and future dilemmas and actions take place accordingly.

As the user interacts with the system a record is kept, in the third person, of everything which occurs within the storyworld. This story record is available to the user following their experience, as a non-interactive narrative.

IV. Evaluation

Evaluations of both the soap and dinosaur adventure versions of GADIN are discussed in this section. This begins with a discussion of a Turing-style test of the story quality of the soap narratives generated by the non-interactive version of GADIN. This is followed by studies of user experiences in both the interactive soap and dinosaur adventure.

A. Story quality

The Turing Test was originally designed by Alan Turing [22]. In the narrative domain it can be interpreted as requiring a human to be unable to determine whether a narrative was created by a human or a computer. In this test two narratives are given to a human reader, one written by a human, the other by a computer. The reader must then decide which they think was written by a computer.

The human-authored narrative used in this experiment has been taken from a television soap opera. This was chosen as randomly as possible, although there were certain restrictions due to limitations of the system. The two restrictions were: there could be no familial relationships in the selected narrative as these had not yet been implemented in the system; and there should be a strictly bounded subset of characters in the chosen narrative.

1) If the narrative is finite check to see if the storygoal has been achieved. If so, end the narrative. If not, continue checking this after each stage in this process.
2) Let the user act if they did not act last, and respond appropriately to the user (with a triggered dilemma or utility-based response).
3) If there is a storyplan present the actions (or dilemma) at the next level of it if possible.
4) If this was not possible:
   a) Check to see if there is a plan for a dilemma to the user awaiting presentation and if so present as much of it as possible, depending on (and integrating) user actions.
   b) Let the user act if they did not act last, and respond appropriately to the user (with a triggered dilemma or utility-based response).
   c) If there was no such plan available or presentation failed:
      i) Let the user act if they did not act last, and respond appropriately to the user (with a triggered dilemma or utility-based response).
      ii) If the restrictions on dilemma experience ordering do not prevent it, check to see if there is a plan for a dilemma indirectly involving the user awaiting presentation and if so present as much of it as possible, depending on (and integrating) user actions.
      iii) Let the user act if they did not act last, and respond appropriately to the user (with a triggered dilemma or utility-based response).
      iv) If there was no such plan available or presentation failed:
         A) If the restrictions on dilemma experience ordering do not prevent it, check to see if there is a plan for a dilemma not involving the user awaiting presentation and if so present as much of it as possible, depending on (and integrating) user actions.
         B) Let the user act if they did not act last, and respond appropriately to the user (with a triggered dilemma or utility-based response).
         C) If there was no such plan available or presentation failed:
            - If the narrative history does not prevent it, check to see if there is an event available for presentation and if so present one.
            - Let the user act if they did not act last, and respond appropriately to the user (with a triggered dilemma or utility-based response).
            - If it is a finite narrative and the storyplan has been deemed to have failed then re-plan for the same or a new storygoal.
            - If the user has still not been presented with anything and a long period of time has passed then select and present a random action.

Fig. 9. An overview of the GADIN narrative generation process.

The human-authored narrative has been transformed to match GADIN's output and level of abstraction. This does not invalidate the Turing test since the intention is to evaluate the quality of the essence of the narrative and not the manner in which it is presented. The comparison is of the structure of the narrative.

The GADIN-authored narrative used in this experiment was not selected but was the first to be generated from the given storyworld state. It was ended when at the same length as the television soap narrative. Although the GADIN system (like the television soap) would continue infinitely from this point the later narrative content was not relevant for purposes of this comparison.

To ensure a fair comparison it was necessary that both narratives began with the same storyworld state. Thus the GADIN system was given an initial storyworld equivalent to that at the start of the selected television soap storyline. The characters were limited to only those included in the selected story, but the actions and dilemmas included all of those available to the GADIN system.

It was important to ensure that subjects who were familiar with the television soap narrative selected did not answer the survey. The names were changed to be the same in both stories, to anonymous names. This meant that any bias due to familiarity with the soap, if not the narrative in question, would be removed. A narrative which had taken place a year previously to the evaluative comparison was selected to further reduce the chance of familiarity due to casual observation.

To make the narratives more readable they were adapted slightly. The actions of a character moving between locations were removed in both versions – this is something which is relevant in the experience but not in the subsequent telling. For the television soap it was necessary to transcribe the events in the form output by GADIN. Using this style of writing is an obvious limitation of the narrative quality, but in this evaluation only the core components of the narrative and its structure are required to be compared and thus this is sufficient. In both versions it was essential not to impose any reasoning on the characters in the wording. In the transcriptions of the television narrative the characters' actions when not part of the main storyline are excluded as these are not relevant for the purposes of this comparison of a single narrative involving

one set of characters.

The question asked of the participants, which includes the two story transcriptions, is shown here:

If you recognise one of these stories, or think you might, from a soap then please do not continue as you will bias our results. One of the following two stories is from a real soap opera and the other was generated by a computer.

Please read the two stories, and decide which you think was written by a computer.

STORY 1: Jane and Tom are in a relationship. Jane becomes pregnant. She decides not to keep the baby. Tom proposes but Jane rejects him, ending the relationship. Nick and Sally go to the shop. Nick flirts with Tom. Sally flirts with Tom. Tom must choose whether to partner Sally or Nick. He decides to go out with Sally. Nick no longer fancies Tom. Nick starts to go out with Rich. Sally wants to start an affair with Rich but he chooses not to. Tom expresses his disapproval of Rich's relationship with Nick, but Rich decides to continue the relationship anyway. Jane and Sally stop liking one another. Tom starts an affair with Jane.

STORY 2: Tom and Jane are in a relationship. Rich and Sally flirt with each other, and start going out. Rich wants to start an affair with Tom but he chooses not to. Rich splits up with Sally. Tom stops liking Rich. Rich and Sally start going out with each other again. Rich ends the relationship with Sally. Rich starts a relationship with Nick. Tom expresses his disapproval of this relationship but Rich ignores him. Rich and Tom start having an affair. Tom proposes to Jane and she accepts. Rich tells Jane about his affair with Tom. Jane ends her relationship with Tom. Tom and Rich start a relationship.

Which story do you think was written by a computer?[1]

The survey was divided into two groups of participants. One consisted of those who regularly view soaps (although not the soap in question). These participants were targeted through posts on 4 English soap forums [23]–[26]. Given the focus on English soaps throughout it was important that these forums were used only by an English audience, as the style and content of soaps does vary to some extent between nationalities. The second group contained those who regularly play computer games. For this the survey was posted on 2 English games forums and 1 international [27]–[29]. To ensure that those familiar with the television narrative did not answer the survey it was additionally requested in the forums that anyone who felt that either story was known to them did not continue.

The results were as follows:

Participant type	*Number*	*Correctly identified stories (%)*
Soap viewer	42	24 (57.1)
Game player	85	49 (57.6)
Total	127	73 (57.5)

Under the null hypothesis, participants are guessing randomly which story is computer generated. However, with a one-tailed alternative hypothesis that they are more likely to discern that story 1 is computer generated, then from a Binomial test, this sample is significantly different from random guessing ($z = 0.055$, $p < 0.05$). With a two-tailed test, there is no significant difference ($p = 0.110$) suggesting that whilst there is an effect, it is a quite small one – which is probably due to the ordering bias in the narratives.

Although only one story has been used for this comparison this gives a strong indication that it will apply to other stories as well, particularly since the story was not specificially selected. This suggests that the GADIN system is capable of generating a structurally sound narrative which is not discernibly different from the structure of a television soap opera.

B. Interactive soap evaluation

Most of the requirements of both computer games and narrative theory will apply to interactive narrative, as will additional considerations unique to interactive narratives. Laurel's 1986 thesis (see [10] for a summary and continuation of this work) was the first to address the concept, and to identify the main such requirements, although the ideologies were never implemented. Subsequent work [1]–[8], [11], [13]–[16] has considered these in greater depth and breadth, and has resulted in the production of a range of interactive narrative systems. The criteria considered in previous research provide the basis for those measures for the success of an interactive narrative system considered in this evaluation. This involved a number of users answering a questionnaire following an experience with the soap version of the GADIN system (section III-A). The questionnaire evaluted the following main areas of the users' experiences: the dramatic interest of the narratives; how immersed they were in the experience; whether they were free to act as and when they wanted to; whether they felt their actions were having a long and short term effect; and the likelihood of their replaying.

There were a total of 47 users selected for this evaluation.[2] Of these 12 experienced a version in which there were no dilemmas, and thus no planning, but still utility-based responses and random character actions. The averaged results from these experiences are discussed in this section. In this discussion users are considered to be 'regular' viewers of soaps or players of games if they stated that they do so more than once a fortnight. All of the averages are given as means with associated standard error.

[1]Story 1 was written by the GADIN system. Story 2 is from Hollyoaks, with the original characters changing names with the following mappings: Craig Dean → Tom ; Sarah Barnes → Jane ; John-Paul McQueen → Rich ; Hannah Ashworth → Sally ; Spike → Nick

[2]Some users only entered the system and did not experience a narrative. To remove the bias which would be caused by results based on experiences which did not take place those experiences which lasted less than 5 minutes have been excluded. In addition the experiences for which the questionnaire was only partially completed were not included.

Users were asked to spend at least 5 minutes experiencing the narrative before they would be eligible for the offered prize. On average users of the version with dilemmas played for 19±2 minutes, significantly longer than they were asked to play for. Since experiences which were shorter than 5 minutes, or in which the questionnaire was not completed, were disregarded there is some bias in this. Users of the experience without dilemmas played for a shorter time, on average 14±1.5 minutes, again excluding shorter experiences and those with an incomplete questionnaire. Those users of the version with dilemmas who regularly both watch soaps and play computer games played for significantly longer than any other group, over 25±5 minutes on average. This is approximately the length of the utilised soap opera episodes.

In the questionnaire only three of the criteria scored below 4 out of 7 on average, as discussed in this section. This is very positive given that this was achieved despite the limitations on the system – in the number of actions and dilemmas available and the interface restrictions. For some of the criteria these ratings were higher for particular groups of users, as discussed. This discussion is divided according to the questions presented to the users. This is accompanied by a summary of the mean ratings given in answer to each of the questions, each with associated standard error, in the order:

- All users of the experience with dilemmas (*all*);
- Users of the experience with dilemmas who regularly play games (*games*);
- Users of the experience with dilemmas who regularly watch soaps (*soaps*);
- Users of the experience with dilemmas who regularly both play games and watch soaps (*both*);
- Users of the experience without dilemmas (*without*).

Were your options always clear to you?

all	*games*	*soaps*	*both*	*without*
4.1±0.2	4.4±0.3	4.0±0.3	4.4±0.5	4.0±0.5

Was what was happening in the storyworld always clear to you?

4.2±0.3	4.5±0.4	4.3±0.4	4.5±0.5	4.3±0.5

Users felt that their options and the happenings in the storyworld were generally clear. This is similar for all categories of user. The remaining results can thus be discussed with the knowledge that the users were evaluating an experience which was reasonably clear for them.

How interesting would you say the story was?

all	*games*	*soaps*	*both*	*without*
3.5±0.3	3.8±0.4	3.8±0.4	4.5±0.7	3.3±0.4

How dramatic would you say the story was?

3.5±0.2	3.8±0.3	3.7±0.3	3.8±0.4	3.1±0.4

The mean ratings for how interesting and dramatic the experience was were below average. However for regular soap viewers and regular game players these were slightly higher. Those who both regularly watch soaps and regularly play computer games (the expected audience of such an experience) considered the interestingness of the narrative to be higher still. In the version in which the users did not experience any dilemmas the interestingness of the experience was rated lower, with a similarly lower rating for how dramatic they felt their experience was.

How immersed were you in the experience?

all	*games*	*soaps*	*both*	*without*
3.0±0.2	3.3±0.3	3.0±0.3	3.0±0.7	3.3±0.3

Users generally experienced a below average level of immersion. This mean was the same for all categories of user. Many users made comments implying that the interface had a negative effect on this. For example one user commented that they did not become immersed as "the interface meant that I was quite aware that I was conducting an experiment". Without dilemmas a similar level of immersion was achieved.

Did you feel that the storyworld was plausible?

all	*games*	*soaps*	*both*	*without*
4.3±0.3	4.4±0.4	4.1±0.4	4.5±0.4	4.1±0.5

Did you feel that the other characters were plausible?

4.1±0.3	4.3±0.4	3.9±0.3	4.3±0.5	4.2±0.4

The plausibility of the storyworld was rated around average, as was the plausibility of the characters. The plausibility of both the storyworld and the characters were rated similarly by users of the experience without dilemmas.

Were you able to act whenever you wanted to in the story?

all	*games*	*soaps*	*both*	*without*
5.1±0.3	5.3±0.3	5.1±0.4	4.8±0.5	4.4±0.5

Were you able to act in the way you would expect to be able to?

3.9±0.2	4.1±0.3	4.2±0.3	4.1±0.5	3.4±0.5

The users felt that they were able to act when they wanted to, which may suggest that the interactive nature of the system has been successful. They also felt that they were able to act as they wanted. These were slightly lower for the version without dilemmas but not significantly, except perhaps to suggest that users felt more need to select actions when there was less narrative structure.

Did you feel that your actions were having as much immediate effect as in the real world?

all	*games*	*soaps*	*both*	*without*
4.2±0.3	5.0±0.3	3.6±0.4	4.3±0.4	3.5±0.4

Did you feel that your actions were having as much long-term effect as in the real world?

3.9±0.3	4.4±0.3	3.1±0.3	3.6±0.5	2.8±0.3

Users found that there was a fairly clear immediate and long-term effect of their actions. Regular game players rated both of these with a higher mean. As these users will be more familiar with computer games this suggests that the effects are in accordance with their expectations in that medium.

In the experience without dilemmas users felt that there was less effect of their actions. They rated the immediate effect of their actions slightly lower. The greater difference is in the long-term effect of their actions, without dilemmas users rated this much lower. This may imply that the dilemmas allow the user to see a long-term effect of their actions on the overall narrative.

How hard would you say major decisions were to make?

all	*games*	*soaps*	*both*	*without*
4.7±0.3	4.8±0.4	5.1±0.3	5.1±0.5	4.3±0.4

Users felt that major decisions were fairly hard to make. Since users always took between 10 and 30 seconds to make their decisions on dilemmas it appears that a reasonable amount of thought was put into such decisions and that these were fairly hard for them to make. Those users who did not experience dilemmas still rated the difficulty of major decisions quite highly. This is probably due to the lack of direct relation between this question and the dilemmas – but users never took as much time to make other decisions.

How likely would you be to play again?

all	*games*	*soaps*	*both*	*without*
4.1±0.3	4.9±0.4	3.9±0.4	5.0±0.6	3.5±0.4

If this was not your first GADIN experience, how different was the narrative this time?

Users generally wanted to repeat their experience, particularly regular computer game players. This was slightly lower in the experience without dilemmas. Most of the users were experiencing the system for the first time. However the 3 who were not all felt that their experience was significantly different from the last (giving this question ratings of 4, 6 and 6 – all out of 7).

Overall how would you rate your experience?

all	*games*	*soaps*	*both*	*without*
3.9±0.2	3.9±0.3	3.7±0.3	4.0±0.5	3.8±0.4

The overall experience rating was the same for both the experience with dilemmas and that without. This indicates that users enjoyed interacting with the soap world, regardless of the dilemmas. However the users of the experience without dilemmas gave generally lower scores for the remaining criteria. For users of the experience with dilemmas the mean of these ratings is the same as that for all of the other scores, implying that these criteria accurately reflect the user's opinion of their overall experience.

C. Dinosaur adventure evaluation

The dinosaur adventure version of GADIN is detailed in section III-B. In this section an evaluation of that version is discussed. The target audience of this version is children between the ages of around 7 and 10. Younger children would need someone to read for them and older children are less likely to enjoy the world. This experience was provided to 6- to 9-year-old pupils at a first school as one of their available activities during a school day.

A questionnaire was used following each user's experience. It was decided to ask the questions verbally and transcribe the answers, as this is a communication method which the children are more comfortable with. They were still reluctant to give much information and it was necessary to limit the number and scope of questions as a result of this. Although some conclusions can be drawn from the questionnaire, as discussed here, it is necessary to supplement this with observations of the users and their responses to the experience.

The children appeared to really enjoy their experiences with GADIN. They wanted to continue playing for as long as possible, with one child using his lunchbreak to play, although another unfortunately had to leave as it was hometime. At one stage there were a large number of children gathered around the interactive whiteboard asking if their turn could be next. All of the users' comments on the story were positive, which may suggest a high level of interestingness.

When they started their experience the users had some questions about what they had to do, but they soon realised that it was entirely up to them and understood the interface well enough to be left to interact. They appeared to find the interface easy to use, indeed one child said "this is actually quite easy" after the first action. The only questions later asked were when presented with dilemmas, with one user asking "what do I do?", greatly perplexed. This suggests that the dilemmas did present difficult decisions for the users. It is believed that not enough dilemmas were experienced and this needs to be improved, possibly by planning for dilemmas or by incorporation of additional dilemmas in general form.

One of the major expressed disadvantages of the interface was not everything being interactive, in particular the pictures. The users wanted to be able to click on the pictures and observe a response. This may suggest that it would be an improvement if the user were able to choose their actions from a combination of graphics and words as opposed to purely textually. This would require the interface to be redesigned appropriately. There was also some confusion when there were references in the text, for example to dinosaur nests, for which no accompanying picture was shown.

The level and volume of reading was right for the older users, but the younger (6-year-olds) found it difficult, and with one user it was necessary to read aloud the longer sections of text – otherwise this would have had an adverse effect on their experience. The dinosaur names were confusing for those children who knew nothing about dinosaurs.

The majority of the users expressed emotions regarding the characters and the storyworld which suggested that they found these plausible. For example one referred to the dinosaurs as being "mean". This may have been influenced by their previous knowledge of the dinosaurs rather than the experience but many did not have any such prior knowledge. Observation of the users showed that even when they had to wait to be able to select an action they remained interested in the system and did not get distracted by the other activities going on around them. As soon as they were able to do so they started to select their next action.

Over half of the children felt that they changed nothing through their presence in the storyworld, mainly the younger users (6 and 7 year olds). This is interesting since they were present and thus their actions were having an effect, even if only to make dinosaurs like them. That they could not see this perhaps suggests that they see stories as the same no matter what and did not understand the effect of their actions. This may suggest that interactive narratives may be inappropriate for the younger end of the evaluated age group as they will be more acceptant of being directed towards particular actions and expect this.

The children tended to feel that they could not always do what they wanted. The reasons for this varied greatly but were

generally ambitious and required for example "to be there in real life". None of the expressed reasons for not being able to do what they wanted were related to the actions which were available to them within the world, except where they were restrained by consistency, such as not being able to steal an egg when they had not found a nest.

There were problems with the storygoal. Since it is communicated before the narrative begins some of the children did not understand it due to a lack of familiarity with the characters. Others quickly forgot it, with one later expressing frustration as she "wanted to get back to the beginning" and find out what the storygoal was. The 9-year-old users generally understood and remembered the storygoal and appeared to be pleased when they were able to achieve it. This suggests that this method of communication of the ending is inappropriate for younger children.

Although nearly half of the users said that they did not want to play again this is thought to be due to a lack of understanding of the question. Some users who said they would not want to play again returned later, when others were interacting with it, to ask if they could do so again. It is believed that they understood the question as requiring them to play again straight away and they did not want to since the evaluation was being performed at the same time as other activities, and the children had art projects to finish. One user enjoyed the experience so much that he asked for a copy of the system.

V. Discussion

In this section the results of the evaluation are discussed. In addition the scalability and transferability of interactive narrative systems, including GADIN, are discussed in more detail.

A. Evaluation results

This section discusses the evaluation in accordance with the following desirable features: interestingness, immersion, user freedom, clear effect of user actions and replayability of the system.

The interestingness of a GADIN narrative is dependent on the dramatic interest of the dilemmas which are defined, in general form, for the domain. That the dilemmas are interesting is suggested by the comparison of experiences with and without dilemmas in which users found the experience with dilemmas to be more interesting, and that users felt that the dilemmas required making difficult decisions. The interestingness will also depend on the appropriateness of the narrative to the audience. It has been shown that the target audiences of the existing versions of GADIN find these narratives interesting (section IV-B and IV-C). The story quality evaluation discussed in section IV-A further implies that the interestingness of these will be equivalent to the interestingness of soaps.

The immersion felt by users of the GADIN system is below average, which can thus not be considered to have been satisfied. This is believed to be at least partially due to the interface restrictions and waiting times. Plausibility of characters and the storyworld are a significant factor in user immersion and the questionnaire shows that users felt that the characters were reasonably plausible, as was the storyworld. The users of the dinosaur adventure appeared to be very immersed in the experience. They were not distracted even when they had to wait and despite the surrounding potential for distractions.

Users were generally able to act as they expected to, and did not seem to find the lack of ability to select every conceivable action a significant restriction, but only expressed that ideally they would be able to. Since the system is currently control-based users were always able to act when they would expect to be able to, and the questionnaire showed that they felt this to be the case. Users of the dinosaur adventure version of GADIN seemed to feel a lack of freedom but the alternative actions suggested were inconsistent with the narrative. The level of freedom experienced by users of the GADIN system is reasonable but not as high as would be preferable or as is achievable.

The questionnaire shows that users felt their actions to be having both an immediate and a long-term effect on the narrative. The long-term effect was felt to be lower in the experience without dilemmas, which may suggest that the incorporation of dilemmas gives the user a clearer long-term effect of their actions. The older users of the dinosaur adventure version of the system felt that they were changing the narrative through their actions.

Users generally felt that they would like to experience a GADIN narrative again. In the soap version this was particularly the case for regular computer game users who are the most likely to enjoy (and be familiar with) this type of narrative. Those who experienced the GADIN system more than once generally felt that the narrative on subsequent experiences was significantly different.

B. Scalability

Traditional narratives vary in length depending on the medium in which they are told. A book may take days to complete, and will generally be read in installments. It may have sequels which will make the overall narrative last even longer. Plays may last 2 or 3 hours. According to Field [30] the average (and ideal) length of a film is 2 hours. There also exist short versions of these storytelling methods. These include: short stories; short films [31]; and short plays or sketches. Television serials and soaps continue over a long, potentially infinite, period of time.

Extending the length of interactive narratives will enable the user to become more involved in the narrative over a prolonged period of time. Current computer games involving a story may take days, weeks or even months to complete. More extensive research (with successful systems) would be required to determine the ideal length for an interactive narrative. This would result in the user being fully involved in the story and genuinely caring about the outcome and the narrative as well as the interaction.

The use of a plot graph by the Oz Project [1] means that there is a large amount of pre-definition required for even short

narratives. Experiences with the Façade [3] system and the associated narratives – although much longer – also centralise essentially on a plot graph - with far more plot points at a much lower level than it is possible to explicitly draw. It took 2 man-years to create a 15-minute experience [3].

The length of Erasmatron [2] narratives will depend on the path the user takes through the storyworlds, with an upper limit being imposed by the size of the world. The DEFACTO [4] and IDtension [6] stories both require the creation of new content, with ordering, to generate longer stories. In contrast, the GADIN system is able to generate narrative content from generalised forms of dilemmas, and does not involve any ordering constraints. The systems are potentially scalable to longer narratives but would need a large number of changes for this to be possible.

The Mimesis [8] architecture is intended to be able to generate interactive narratives with any game engine and thus of any length. The techniques used by the I-Storytelling [5] group do not lend themselves to a longer narrative due to the large amount of content and ordering pre-definition required.

The IDA [7] experience has a fixed storyline. This means that both this story and its user model, which would be very complex with even a short narrative, must be pre-defined. The INTALE [15] techniques require all possible endings to be pre-defined, meaning that a lot of initial input is required to still enable the user to have freedom in the later stages of the narrative. The FearNot! [16] system also requires a large amount of content and ordering pre-definition for longer narratives.

The OPIATE [13] and PaSSAGE [14] experiences are more similar to computer games and will thus last longer than most interactive narratives. However the essential story structure of OPIATE is only as long as a Russian folk tale, and it is likely that if the experiences last longer the dramatic interest will be reduced. Scalability to a longer version of the PaSSAGE narratives is considered and certainly appears to be possible.

As discussed in section V-A GADIN narratives have been demonstrated to have a reasonably high level of interestingness. The techniques discussed in section II-K ensure that the GADIN system is capable of generating a narrative of infinite length. The method of relaxing restrictions on applicability leads to reachability between states, meaning that from any storyworld state there will be another state possible and thus the narrative – and presentation of dilemmas – can continue. Although the state information is finite and thus for the narrative to continue infinitely there will be repeated dilemmas, the context of and prior history to these will be different each time. Dilemmas are used by GADIN to provide dramatic interest in the narrative. The continuation of these dilemmas indefinitely within the narrative means that the dramatic interest will be sustained.

C. Domain independence

Domain independence is an essential criteria for interactive narratives. It is not sufficient to have a system which is successful in a single domain. This will appeal to a restricted audience and will greatly limit the potential for narratives produced. It must be demonstrated that the techniques utilised in the design of the interactive narrative system will apply to a range of domains by creating entirely separate storyworlds. A range of narratives must be generated within each of these worlds.

The basic technique of a plot graph structure is applicable to any domain, as has been shown by the Oz Project's work providing the basis for Façade. The application of these techniques to new domains requires a large amount of pre-definition, essentially consisting of re-writing the plot graph for the domain.

The DEFACTO system has only been implemented in a single storyworld and does not give consideration to further applicability. Although the basic techniques used in the IDA system should apply in any domain pre-authoring the entire narrative is a significant constraint on this.

The Erasmatron and Mimesis systems are designed as general architectures which will apply to any domain. There have been a number of storyworlds created for the Erasmatron system – which is aided by the story development tool.

PaSSAGE encounters would need to be defined for each domain, and in less computer game-oriented domains the player model is less likely to be applicable, as it is based specifically on computer game players. If functions such as Propp's were available for domains other than Russian fairy tales then it would be possible to apply the techniques used in the OPIATE system to these domains. The techniques used in INTALE are potentially applicable to any training scenario, but as the system is not designed for generality fundamental re-writing would be necessary. The FearNot! system is designed specifically to address bullying scenarios but has applicability to other role-playing scenarios, particularly in the classroom.

The IDtension, I-Storytelling and Façade systems have been, or are currently being, applied to more than one domain. For each new domain significant content creation is likely to be required in all of these systems.

The exploitation of the high level of reusability and generality in the soap storylines means that the soap version of the GADIN system can be easily applied to any soap of this type, simply by adapting the involved characters (and their initial feelings) appropriately. This demonstrates a certain independence as the system is applicable to any existing soap of this type (and the storylines will be appropriate), to a new soap or to the user's self-defined soap world.

The application of the GADIN system to a generic children's story domain demonstrates that the techniques are applicable to children's adventure stories.

Soap stories are very relationship-centric. The audience will become involved due to the relationships between the characters, how they interact with one another and their changing feelings for each other. This is not true of the dinosaur domain, but the techniques used by GADIN are still applicable within this domain. Ensuring that a narrative ends in a timely and coherent manner is a very different challenge to maintaining the dramatic interest of the narrative for an infinite length of time. That the techniques used by this system can be applied to both types of storytelling demonstrates the robustness and fundamental generality of the core of the system.

The fundamental techniques used in the GADIN system are applicable to interactive narrative generation in any domain which makes use of generic storylines of the type discussed in section II-C. The core architecture of the GADIN system is the same, and reusable, for any application domain. Once the domain specific information has been incorporated into GADIN the system will be able to generate narratives within that domain. In creating a new application domain for GADIN it is only necessary to determine and define the storyworld, actions, dilemmas and the necessary components of the required additional features. For the generation of stories GADIN needs only to be provided with the generalised form of dilemmas, with no ordering or sequencing constraints imposed, allowing GADIN to generate unique narratives. The ease of this process is discussed here. This is also relevant to the effort required to extend an existing application of GADIN, for example by adding further dilemmas or actions. This discussion focuses on the ease of implementation rather than design, as this is more quantifiable.

In considering the ease of creating new domains it is assumed that the domain creator has a good knowledge of the GADIN techniques. They will not usually be required to have familiarity with the GADIN code other than those components which are directly relevant to the specification of the new domain, which are discussed in this section. Certain optional features require the domain creator to have a greater knowledge of the code, and details are given whenever this is relevant.

Storyworld: Each character has default traits associated. Which traits are used in the current domain can be easily specified. It is also possible for these to be extended, depending on the domain, but this should not be necessary for most domains. It is fairly easy to extend these traits but does require a slightly greater knowledge of the code. The same is true of the requirements associated with locations and objects. Such extensions will become increasingly less necessary as GADIN is applied to more domains and the knowledge base is extended.

There must be specific instantiations of the characters and their traits, the locations and the objects within the storyworld. This determines the initial state of the storyworld. This can easily be created randomly by GADIN. If the domain creator chooses to specify the initial state of the storyworld then – depending on the size and complexity of the storyworld – this can require a large amount of time to complete. However it is simple to do and follows a clear structure, the time is only required since entering names and traits for a large number of characters inherently takes time. Each character, location and object which may appear in the narrative must have all of the traits or requirements for the domain individually instantiated. Alternatively the domain creator could choose to make it contingent on the user to provide this information.

Actions: Adding actions to the GADIN knowledge base is relatively straightforward. It follows a fixed and general structure. There is more complexity if the preconditions or effects require the creation of new categories of propositions. As the knowledge base grows the addition of new propositions will become less necessary.

In applying GADIN to a new domain many actions from previous applications are likely to be reusable, for example the action of moving between locations. Once the relevant actions have been identified it is simple to specify which actions will and will not be appropriate for the current domain. There are also many actions which will have core similarities, such as any characters embarking on a mutual relationship (for example a partnership or friendship). These actions are defined in general form in GADIN, and only the relationship (and output form) need to be changed for each new action of this type.

The applicability of actions is likely to vary depending on the domain. It is necessary for the domain creator to specify all of the possible applicability conditions for each new action which is added to the knowledge base. This can be complex, depending on the number of possibilities for each action.

Dilemmas: The greatest difficulty in incorporating dilemmas into GADIN is in their identification. Following this each dilemma should be generalised and can then be easily instantiated – utilising the clear and general structure which GADIN uses for dilemmas of that category. The specifics of each dilemma are mainly dependent on their category, and are thus already present. There are additionally many dilemmas which will be reusable between domains.

Events: If the incorporation of events is appropriate for the new domain then these need to be identified. They can then be made available to GADIN, following the architecture's clear general structure for events. It is straightforward to add a new event. Some events will be reusable between domains. For example a party will be applicable to many domains, and although the motivations may differ this is easy to adjust without changing anything further regarding that event.

Knowledge: In some domains the use of knowledge is very simple to implement, for example in the dinosaur adventure version of GADIN the information available to the user is restricted by their current location. When the domain requires that there is a more intricate representation of knowledge, as discussed in section II-I, it will be necessary for the domain creator to decide what will be represented in this manner. Once this has been determined the knowledge items must be added to the list of possibilities. Each knowledge item must then have associated: preconditions; reasons to share it; and when it will become (un)known by characters. Initial knowledge held by characters must also be specified.

The user model: If a user model is to be incorporated in the narrative generation process the aspects of the user which should be modelled must be identified by the domain creator. Each dilemma and action should then have associated the aspects it updates. Each dilemma will also have associated the balance of criteria involved in predicting its outcome. Once the user model aspects and appropriate updates have been determined this is straightforward to implement.

Infinite or finite narrative: It must be determined whether a finite or infinite narrative will be generated. If the narrative is infinite then there are no further requirements for the domain creator. If the narrative is finite the domain creator may add restrictions on the possibilities for the storygoal. This is optional (depending on the domain) and requires greater

knowledge of the GADIN code.

VI. Outlook

The GADIN system is capable of sustaining the dramatic interest of generated narratives over as long a time period as is required for the domain (section V-B). The core architecture and techniques are easily applicable to new domains (section V-C). This section discusses some of the limitations of the GADIN system, and proposes solutions.

As more characters and actions are included the time spent planning becomes unreasonably long, thus reducing the extendibility of the GADIN system within a given domain. The time increases exponentially with the number of characters and the number of actions. The number of locations does not adversely affect the planning speed. This is due to the large number of plans which must be found for every combination of characters and every dilemma. Since a GraphPlan algorithm is used and the planning graph is used for multiple dilemmas increasing the number of dilemmas does not have the same impact on the speed.

It is possible that a different planning algorithm would improve the speed. Barros and Musse [32] analyse potential planners for interactive narratives, but the adjustments made to ensure specificity of the planning algorithm to the domain means that changing the planner is unlikely to resolve this issue. The problem is not the time taken to find a single plan but the volume of plans that must be found.

A potential solution would be the use of a form of hierarchical planning in which there is default knowledge of possible plans, with variables, which lead to certain propositions. This could be generated for each new storyworld and stored so that it could be referenced whenever a narrative was generated in that storyworld.

Another possibility would be for the actions currently available in the GADIN system to be considered as directions rather than the full actions. Storyworld characters would then improvise to carry out these directions, in a manner dependent on their personality. They could create low-level plans to achieve the effects of the required action.

The GADIN system aims to create the dramatically interesting components of a narrative not the visualisation of these. The elements of the narrative are abstracted to some extent but the essential elements and the components of dramatic interest remain. However from the analysis in accordance with the compellingness criteria it is apparent that immersion and user freedom must be improved for this to be a fully compelling interactive narrative system. The questionnaire implies that this limitation is mainly due to the interface.

Although a graphical interface is likely to complicate the further applicability and extendibility of the GADIN system it should increase the sense of immersion and plausibility to the user of the characters and the storyworld. It may be possible to utilise an existing graphical engine for GADIN, for example Second Life. With a graphical interface it may also be possible for the dilemmas to become less explicit. In the original literature they are an integrated component of the narrative but in GADIN are explicitly presented to the user. It would be advantageous if the dilemmas could be integrated more subtly into the narrative experience while maintaining the dramatic interest.

Appendix A
An example soap story

An extract of the narrative created in one user's (jim) experience with the soap version of GADIN is shown here. Line numbers are listed for reference purposes. User input is shown in bold, with a **n** used to signify the user returning control to the system. jim is male and young. In the initial state the other characters are tess, edna, mary and fred. The initial feelings are as follows:

- tess likes edna, mary and fred and fancies edna and jim.
- edna likes tess and fred and fancies tess.
- mary neither likes or fancies any other characters.
- fred likes tess and mary but fancies noone.
- jim (the user) likes mary and fancies edna.

1 -**n**
2 Action is mary moves between house and club
3 -**move club**
4 You are already at the club!
5 -**n**
6 Action is edna starts to fancy you
7 Action is mary gets drunk
8 -**n**
9 mary offers to buy you a drink. Will you accept?
10 **y**
11 You accept the drink from mary and get drunk
12 -**n**
13 edna asks you out.
14 Would you like to go out with edna
15 who you fancy and who fancies you?
16 **y**
17 You are now partners with edna
18 Action is mary starts to fancy you
19 -**fancy mary**
20 You start to fancy mary
21 You are going out with edna.
22 Would you like to cheat on edna with mary
23 who you fancy and who fancies you?
24 **y**
25 You are now cheating on your partner edna
26 with mary.
27 **n**
28 Action is fred encourages edna
29 to break her principle of not stealing
30 fred is encouraging edna to betray
31 her principle of not stealing
32 edna likes fred and decides to agree
33 to betray her principle
34 edna no longer holds the principle
35 of not stealing
36 As a result of this choice, fred likes edna more.
37 -**flirt tess**
38 You flirt with tess

39 **-n**
40 Action is tess moves between
41 park and club
42 Action is tess flirts with you
43 **-n**
44 You are going out with edna,
45 and cheating on edna with mary.
46 Would you like to leave your partner edna
47 to go out with mary?
48 **y**
49 You are now partners with mary
50 and no longer going out with edna.
51 **-n**
52 You have found out that your friend mary's
53 mother is in an institution
54 Would you like to tell them ('y')
55 or keep this a secret ('n')?
56 **y**
57 You have chosen to share this news
58 Everyone now knows this.
59 **-n**
60 Action is fred moves between house and club
61 Action is fred gets drunk
62 Action is tess gets drunk
63 tess starts to mutually fancy fred
64 tess has to choose whether to partner
65 edna or fred
66 Where tess fancies both potential partners
67 and the feeling is mutual
68 tess decides to go out with fred
69 They are now partners,
70 and edna no longer fancies tess
71 **-disapprove tess fred**
72 You express your disapproval of
73 the relationship between tess and fred
74 **-n**
75 You disapprove of tess's relationship with fred
76 tess ignores your opinion as
77 she does not like you.
78 **-drunk**
79 You get drunk
80 **-n**
81 Action is fred gets drunk
82 **-n**
83 Action is fred starts to fancy you
84 **-n**
85 Action is edna starts to fancy you
86 You are going out with mary.
87 Would you like to cheat on mary with edna
88 who you fancy and who fancies you?
89 **n**
90 You have chosen not to cheat on your partner.
91 As a result of this choice, edna fancies you less.
92 **-n**
93 Action is edna begins to disapprove of
94 fred's relationship due to the age difference
95 edna disapproves of fred's relationship with tess
96 fred likes edna and must thus decide
97 whether or not to respect her opinion
98 and end the relationship
99 fred decides to end their relationship.
100 fred and tess are no longer partners
101 edna now likes fred more.

REFERENCES

[1] J. Bates, "Virtual reality, art, and entertainment," *Presence: The Journal of Teleoperators and Virtual Environments*, vol. 1, 1992.
[2] C. Crawford, *Chris Crawford on Interactive Storytelling*. New Riders, 2004.
[3] M. Mateas and A. Stern, "Façade: An experiment in building a fully-realized interactive drama," *Game Developers Conference, Game Design track*, 2003.
[4] N. M. Sgouros, "Dynamic, user-centered resolution in interactive stories," *International Joint Conference on Artificial Intelligence*, vol. 2, 1997.
[5] M. Cavazza and F. Charles, "Character-based interactive storytelling," *IEEE Intelligent Systems*, vol. 17, 2002.
[6] N. Szilas, "IDtension: A narrative engine for interactive drama," in *1st International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, Darmstadt, Germany, 2003.
[7] B. Magerko, "Story representation and interactive drama," in *1st Artificial Intelligence and Interactive Digital Entertainment Conference*, Marina Del Rey, California, 2005.
[8] R. M. Young, M. O. Riedl, M. Branly, A. Jhala, R. J. Martin, and C. J. Saretto, "An architecture for integrating plan-based behavior generation with interactive game environments," *Journal of Game Development*, vol. 1, 2004.
[9] M. L. Ryan, "Narrative and the split condition of digital textuality," *Dichtung Digital*, vol. 5, 2005.
[10] B. Laurel, *Computers as Theater*. Addison-Wesley Publishing Company, 1993.
[11] J. Murray, *Hamlet on the Holodeck: the future of narrative in cyberspace*. Cambridge, Massachusetts: MIT Press, 1998.
[12] M. Mateas, "An Oz-centric review of interactive drama and believable agents," School of Computer Science, Carnegie Mellon University, Tech. Rep. CMU-CS-97-156, 1997.
[13] C. Fairclough, "Story games and the opiate system," Ph.D. dissertation, University of Dublin - Trinity College, 2004.
[14] D. Thue, V. Bulitko, M. Spetch, and E. Wasylishen, "Interactive storytelling: A player modelling approach," in *AIIDE'07*, Stanford, California, 2007.
[15] M. Riedl and A. Stern, "Believable agents and intelligent story adaptation for interactive storytelling," in *TIDSE'06*, Darmstadt, Germany, 2006.
[16] R. Aylett, S. Louchart, J. Dias, A. Paiva, and M. Vala, "FearNot! - An experiment in emergent narrative," in *IVA*, 2005.
[17] "Fable," Microsoft Game Studios [Computer game], 2004, developed by Lionhead Studios. Designed by Peter Molyneux.
[18] "Grand Theft Auto," ASC Games [Computer game], original version released 1997, developed by DMA Design.
[19] Aristotle, *Poetics*, 350 B.C.E., translation by S. H. Butcher available at http://classics.mit.edu/Aristotle/poetics.html. Last referenced 01/10/08.
[20] A. Blum and M. Furst, "Fast planning through planning graph analysis," *Artificial Intelligence*, vol. 90, 1997.
[21] A. Shepard, *The Business of Writing for Children*. Shepard Publications, 2000.
[22] A. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, 1950.
[23] http://www.soapforum.co.uk, last referenced 01/10/08.
[24] http://www.digitalspy.co.uk, last referenced 01/10/08.
[25] http://www.soapboards.co.uk, last referenced 01/10/08.
[26] http://www.tvforum.co.uk, last referenced 01/10/08.
[27] http://www.enemydown.co.uk, last referenced 01/10/08.
[28] http://www.na-clan.co.uk, last referenced 01/10/08.
[29] http://www.wcreplays.com, last referenced 01/10/08.
[30] S. Field, *The Screen-writer's Workbook*. New York: Dell Publishing, 1984.
[31] R. Raskin, "Wordless eloquence in COME," *P.O.V: Filmtidsskrift (A Danish journal of film studies)*, vol. 7, 1999.
[32] L. M. Barros and S. R. Musse, "Planning algorithms for interactive storytelling," *Computers in Entertainment*, vol. 5, 2007.