Plan-based Narrative Generation with Coordinated Subplots

Julie Porteous[1], **Fred Charles**[2] and **Marc Cavazza**[3]

Abstract. Despite recent progress in plan-based narrative generation, one major limitation is that systems tend to produce a single plotline whose progression entirely determines the narrative experience. However, for certain narrative genres such as serial dramas and soaps, multiple interleaved subplots are expected by the audience, as this tends to be the norm in real-world, human-authored narratives. Current narrative generation techniques have overlooked this important requirement, something which could improve the perceived quality of generated stories. To this end, we have developed a flexible plan-based approach to multiplot narrative generation, that successfully generates narratives conforming to different subplot profiles, in terms of the number of subplots interleaved and the relative time spent on each presentation. We have identified specific challenges such as: distribution of virtual characters across subplots; length of each subplot presentation; and transitioning between subplots.

In this paper, we overview this approach and describe its operation in a prototype Interactive Storytelling (IS) System set in the serial drama genre. Results of experiments with the system demonstrate its usability. Furthermore, results of a user study highlight the potential of the approach, with clear user preference for presentations that feature interleaved multiple subplots.

1 Introduction

Interactive Storytelling (IS) has emerged as a popular application of AI techniques in particular Planning to new media entertainment, where planning is in charge of preserving story consistency despite variable initial options or in-story interactions (for example work such as [37, 1, 26, 5, 24, 11]). Hence IS represents an interesting new application of planning: one that challenges received wisdom about optimality and knowledge representation, as the shape of plan trajectories determines narrative structure. [21]). This emphasis on story "backbone" consistency has sometimes resulted in departure from the properties of real-world narratives, which far from reflecting a canonical structure, often interleave various sub-plots featuring the same characters, for instance in popular TV dramas, as analyzed by [35]. Nevertheless an enduring challenge remains, which is to improve user acceptance of these automatically generated narratives.

Interestingly, studies of popular dramas and TV series (an important target genre for narrative generation and IS) suggest that individual episodes, even if self-contained, feature multiple sub-plots. For example, the series MASH [29] had double plot-lines per episode whilst more recent series such as SEINFELD [31] and ER [27] feature even more [33]. Further, the way in which different subplots are interleaved has been shown to be a factor in user acceptance, with individual scenes tending to be short, providing only a slight bit of progression in a given plotline and moving quickly among plots [35].

Previous research has attempted to make narratives most realistic despite the constraints imposed by single-plot single-goal generation. The main methods have sought to improve the presentation of the story at discourse level or to generate narratives with multiple goals. However, none has really explored the generation of coordinated subplots in an effort to reproduce the authoring of real-world drama.

Hence, we were motivated to develop an approach to narrative generation which provided just such a coordination mechanism. Our solution achieves narrative generation and multiplot coordination in a single pass via forward heuristic search through partially ordered subgoals. It represents an extension to our successful approach to single subplot narrative generation which uses authored subgoals to incrementally build up the narrative plan by composition of smaller plan segments, themselves generated using a base planner [22]. We aim to demonstrate that the same landmark model that has been successful to coordinate the pace and trajectories of narrative plans can be extended to represent multiple sub-plots. The coordination of subplots would thus be framed as another high-level landmark ordering problem. This will be managed through heuristic search over the landmark graph, using heuristic knowledge on plot duration, ratios and switching which can be acquired from media studies literature. In this paper we concentrate on narrative generation within our IS system, however we emphasise that a key strength of this incremental forward narrative generation approach is the flexibility it provides to respond to user interaction with the system.

Thus this work addresses an important creativity challenge facing automated narrative generation techniques: to generate output that shares the properties of human-created content but without a heavy authoring burden. Our contribution is the extension of plan-based single subplot narrative generation from action descriptions to the generation of narratives featuring multiple interleaved subplots from a combination of action descriptions and high-level properties of subplots. Clearly these high-level properties – generic filmic knowledge – requires some authoring but we argue that our approach constitutes little additional authoring overhead and greatly increases realism and acceptance (as illustrated by our user study).

Our approach to multiplot narrative generation is fully implemented and integrated within an IS system set in a medical hospital drama (where multiple sub-plots tend to be the rule [35]). This system is used throughout the paper for illustration and was also used in a series of experiments which we report in the paper. Fig.1 provides an overview of the system architecture.

The paper is organised as follows: the next section covers the discussion of related work. The narrative framework and the require-

[1] Teesside University, email: j.porteous@tees.ac.uk
[2] Teesside University, email: f.charles@tees.ac.uk
[3] University of Kent, email: m.o.cavazza@kent.ac.uk

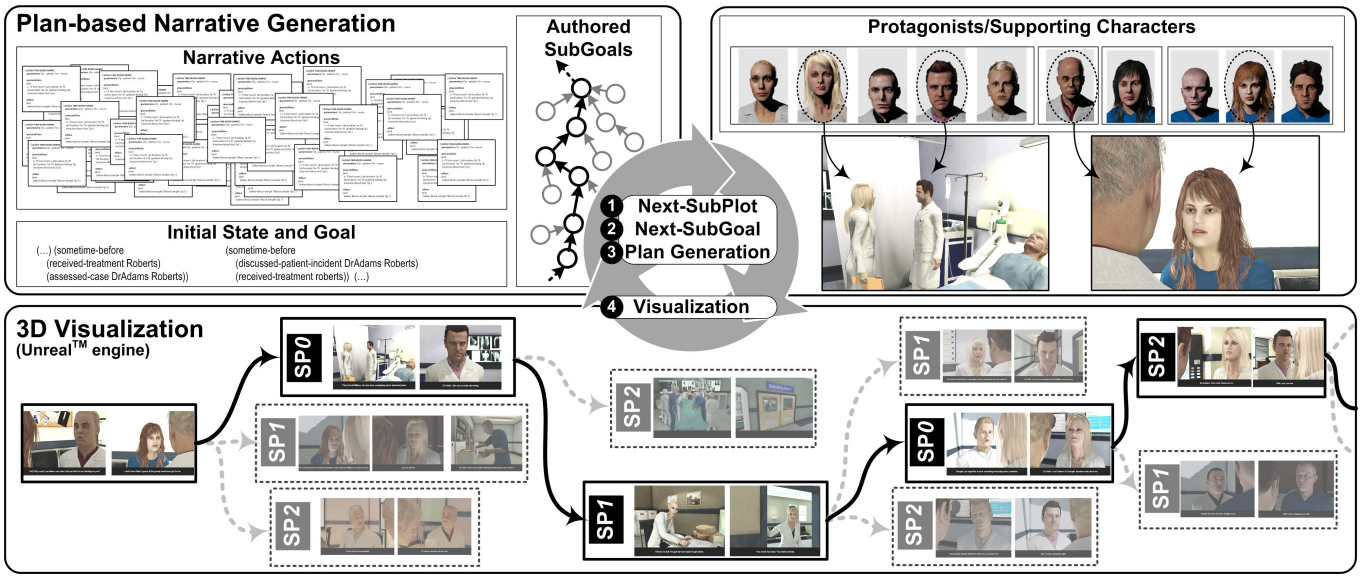

Figure 1. Overview of Multiple Subplot Generation. The system operates in a "plan-execution" loop with narratives generated and visualised incrementally. Each iteration proceeds through the following steps: (1) the selection of the next subplot for the narrative to transition to; (2) for that subplot the selection of the next subgoal to use to generate the next phase of the narrative towards; (3) plan-based generation of the next phase of the narrative using the selected subgoal as goal; (4) as narrative segments are generated they are visualized using the UDK® game engine and shown to the viewer.

ments for multiple subplot narratives are detailed in section 3). Authoring of the narrative domain model is discussed in section 4. The narrative generation procedure is detailed in section 5 and is illustrated with an example (section 6). Results of experiments and user study are reported in sections 7 and 8. In section 9 we conclude and assess the potential of the approach.

2 Related Work

There are many examples of the use of planning for generation of narrative but none to date looking at interleaving of multiple plotlines. A number of these plan-based approaches to narrative generation use authored goals to help shape output narratives and in our work we have drawn inspiration from them. In particular, [23] introduced a notion of author goals and used them to extend a partial order planner in an approach referred to as "complexifying" of the planning process. Subsequently, [21] used authored constraints to help structure narrative as part of a dynamic run-time mechanism that selected constraints for generation of narrative variants. Their approach exploited forward state space planning (essential for supporting interaction), used a standard representation language (PDDL3.0), and provided an author-friendly interface to specify narrative plan dynamics.

The interleaving of multiple subplots can be seen as a discourse level mechanism, for example, used to maintain user interest. Some related research has tackled discourse level aspects by firstly generating narrative and then tackling discourse in a post-processing phase. Examples of this include, [6, 2], for the introduction of suspense, and [14], for cinematic representation. However, our approach generates and coordinates multiplot narratives in a "single pass": something which is arguably more principled and more robust than presenting separate aspects of plot at discourse level.

There are some narrative approaches that have adopted a drama manager approach [15, 36, 18, 32]. Whilst this approach could be used for generation of multiplot narratives we argue that this lacks the generative power and flexibility leveraged by planning along with the author friendly mechanism for declarative specification of narrative structuring control knowledge.

Some work has looked at generation of narratives consisting of multiple quests. For example, [16] use planning for off-line adaptation of authored narratives which include multiple "quests" in the context of game plot lines with focus on plotline adaptation to create new plausible quest sequences. The Crystal Island interactive narrative of [17] featured multiple quest subplots in an approach to user goal recognition aiming to combine multiple plot elements to create rich customized stories in the domain of microbiology with the user playing the role of detective solving a science mystery. [13] presented an approach to the generation of side quests for role-playing games to enhance players sense of agency. We observe that whilst these systems sought to generate multiple quests there was no requirement to interleave them or to conform to the conventions of a particular genre.

Also related are cinematic conventions governing editing and continuity *" ... to tell a story coherently and clearly, to map out the chain of characters' actions in an undistracting way ..."* [4, 34]. One of the advantages of plan-based narrative generation is its power to support narrative causality both at the local and global levels and to ensure coherent character behaviour [38].

3 Coordinated Multiplot Narrative Requirements

Narratives in the domain of television serial drama tend to be built up from a number of interleaved single subplots, where the resulting output multiplot narrative is composed of chains of *segments* each drawn from the individual subplots [33]. Here, as a target for generating such narratives, we identify a core set of requirements for individual subplot and multiplot coordinated narratives. This constitutes generic reusable filmic knowledge, sourced from leading film studies [3, 33, 35] and genre analysis. The rationale for selection of these requirements is given below and are summarised in Fig. 3. These requirements are also summarised in Fig. 2.

SINGLE	Characters	Protagonist with additional characters: allies (support) and antagonists (obstruct).
	Structure	Narrative segments: *introduction* *obstruction* → *resolution* (one or multiple) *exposition* (one or multiple)
MULTIPLE	Characters	Different Protagonist for each subplot Characters can appear across subplots
	Structure	Composed of two or more single subplots Target total for number of segments Target ratio of time on each subplot Subplot segments distributed evenly
	Interleave	Well-formed multiple subplot narrative is: *a sequence of narrative segments with adjacent segments from different subplots*

Figure 2. Summary of requirements for narratives with **single** subplots and **multiple** interleaved subplots. See section 3 for further detail of rationale.

3.1 Individual Subplots

Analyses from Smith [33] and Bordwell [3] provide insights into the roles of characters and structure of plot in individual subplots for our target genre: at the centre of each individual subplot is one character – the protagonist – a goal driven individual who encounters obstacles in the pursuit of their goal; individual subplot structure requires some initial introduction to the protagonist and their goal, followed by one (or multiple) obstructions to them in the pursuit of their goal, leading to some final resolution. Accordingly we require the following outline structure for individual subplots:

- *introduction*: to introduce the protagonist and their goal
- one or multiple occurrences of:
  - *obstruction*: the protagonist encounters some form of obstacle which prevents them from directly achieving their goal
  - *resolution*: either the obstacle is overcome or goal is achieved
- *exposition*: a feature of narratives in this genre is that individual subplots can, at any stage, include phases of "dispersed exposition" [35], where additional information is given to the audience, e.g. plot re-caps or further insight into characters. Hence we allow for subplots to include one or multiple episodes of exposition.

Importantly, this allows for longer length subplot narratives to be built up via composition of multiple phases of obstruction and resolution, interleaved with one or more exposition phases.

Narrative Segments
For clarity, we introduce the notion of narrative "segments" to refer to the different phases of individual subplots that are coordinated to form multiplot narratives – in other words ground instances of *introduction, obstruction, resolution and exposition* as described above.

3.2 Multiple Interleaved Subplots

Multiple subplot narratives are composed from the interleaving of segments from a number of individual subplots. Hence each of these subplots display the properties relating to character and structure described above. However the subplots are not independent because virtual characters can appear in different roles across subplots and the consequences of events in one subplot can impact on the unfolding of the narrative in other subplot(s). As an example consider an episode of ER [27] (episode 15:17) which features one subplot with a protagonist, Dr Carter, struggling in the emergency department after a break, whilst in another subplot with a different protagonist, Dr Banfield who is struggling to conceive, Dr Carter appears as support. Hence we require that each subplot has a unique protagonist but otherwise characters can appear freely across subplots.

Genre analysis of the medical dramas SCRUBS, HOUSE and ER [30, 28, 27] revealed varying numbers of subplots interleaved in episodes, ranging from two upwards. Further, there was no clear pattern of the relative proportion of each subplot: sometimes the subplots were evenly split whereas in others some subplots were more dominant. Hence, to best support these features of human-authored dramas we developed an approach that was flexible with respect to these structural properties. Thus we allow for the relative proportion of different subplots to be specified as part of the input parameters which are used as targets for generation of a given narrative.

Finally, from the genre analysis was observed an even distribution of subplot segments across entire episodes rather than concentrated in different phases which is consistent with observations of rapid quick movement between subplots [35].

Thus rules are required for managing plot switching that can be readily integrated with the overall narrative generative framework in a flexible manner to provide support for all the major features encountered in human-authored dramas. Hence our framework is flexible with respect to structural properties such as the number of subplots to be interleaved, the rate of switching between them, the relative proportion of the narrative per subplot and the overall duration of the narrative. Consequently, we allow for these requirements (summarised Fig. 2) to be specified as part of narrative generator input.

The final requirement of interleaved multiplot narratives is that they are "well formed" which we take to be sequences of narrative segments where *adjacent* segments are from *different* subplots[4]. As an example, a well-formed three subplot narrative (for subplots s0, s1 and s2), could consist of a 12-segment sequence, with 3, 3 and 6 segments from subplots s0, s1 and s2 respectively:

s1 | s2 | s0 | s2 | s1 | s2 | s0 | s2 | s1 | s2 | s0 | s2

Note that, as shown in the system architecture Fig. 1, since narrative generation in the system works in an incremental "plan-execute" loop, a single narrative isn't generated in the traditional sense, rather as narrative plan segments are generated they are visualised to the audience. This incremental approach allows for greater flexibility in the presence of user interaction as detailed in [22] and has been designed to be compatible with various implementations, whether story variability derived from initial parameters (as in the NETWORKING system [20] used in this work) or is due to user interaction [9].

4 Authoring for Coordinated Suplot Generation

The multiplot generation approach we present in this paper extends our plan-based approach to single subplot narrative generation to multiplot narratives. Hence it requires an input domain model consisting of pre- and post-condition actions augmented with authored subgoals which are used to control narrative structure.

For our medical drama domain, the narrative planning actions are those that characterize the genre, such as conflicts over patient diagnosis, treatment, professional rivalries, battles to save patients, romance, domestic conflicts and support for friends and so on. The process of modelling these narrative actions is as detailed in [22, 20].

The approach also requires the domain model to include authored subgoals to provide control over the structure of the narrative as it is

[4] We note this places certain restrictions on the relative target number of segments for different subplots. We assume this is met.

generated. These subgoals can be authored using an intuitive visual user interface (as discussed in [21]) which helps ameliorate some of the authoring overhead. Below we consider further the authoring of the subgoals and their role in coordinated multiplot generation.

4.1 Narrative Segment Subgoals

Our approach is to require that the narrative planning domain model is augmented to include authored narrative subgoals in the style of [23, 22]. These are partial descriptions of interesting states of the narrative world that can be used to generate segments of the narrative. The minimal representational assumption on these subgoals are that they are sets of domain facts.

As an example, consider the narrative segment which introduces a subplot where the protagonist is struggling with pressure of work as shown in Fig. 3. The subgoal here is *shown-pressure-of-work* and the plan generated for this goal ensures that the pressure of work being experienced by Dr Adams (protagonist) has been introduced in the narrative. This example illustrates a form of disjunctive specification of the subgoal (as in [8]). The subgoal is loosely specified, thus making no requirement on how this goal is achieved and thus allowing a range of different plans to be able to achieve the goal depending on the state of the narrative world. The figure illustrates 2 different alternative narrative segments: one generated when the supporting character Dr Miller is happy and a different one when she is angry.

This form of disjunctive specification of subgoals is important since subplots aren't independent and consequences of actions in one subplot can impact on later segments of other subplots. The small example shown in Fig. 3 also illustrates this: suppose if Dr Miller is also the protagonist of a different subplot, then the progress of that subplot can impact depending on whether she has been *harassed-by-colleague* leaving her angry, or happy as a result of *arranged-date*.

4.2 Partial Order over Subgoals

In order to generate such segments, we adopt the approach of [21] and use authored subgoals to represent appropriate narrative situations that can be used to control structure during narrative generation. As illustration, Fig. 4 shows some subgoals and their orders for a single patient treatment subplot, modelled using the PDDL3 modal operators, *sometime-before*, *sometime* and *at-end* (comments show them numbered, sg1-sg4 and goal). For the purposes of generating a single subplot the subgoal *(taken-case DrGreen Roberts)* can be used to generate a narrative segment introducing the doctors' intentions for the patient, whilst the (exposition) subgoal *(discussed-ethics DrGreen DrCook)* generates a segment that gives the audience information about the doctors' views. We note that even this small partially ordered collection of subgoals allows for the generation of a range of differently structured subplots of varying numbers of interleaved segments. For example, the sequence of subgoals $sg1|sg2|sg3|sg4$ could be used to incrementally generate a subplot consisting of 5 segments (narratives generated using predicates such as *(taken-case DrGreen Roberts)* as goals), whereas the sequence $sg5|ex1|sg6$ could be used to generate a very different subplot with 3 segments, including a segment of exposition, for subgoal *(discussed-ethics DrGreen DrCook)*, where insights are given into DrGreen's ethical views.

In addition we also use partial orders over the authored subgoals as a mechanism for structuring narrative content, for example to ensure that the protagonists goals are introduced early in the subplot and that obstacles encountered by the protagonist happen prior to them being overcome. This structuring information can be captured via an intuitive visual interface, as in [21], and then it can be automatically instantiated to ground domain predicates and PDDL3 *sometime* and *sometime-before* modal operators at run time. This provides a much more user friendly mechanism for authoring this information than within the individual domain actions themselves.

(:action patient-consultation
:parameters (?d - doctor ?p - patient ?l - medloc)
:precondition
(and (patient-doctor ?d ?p) (struggle-work-pressure ?d) ...)
:effect (and
(consulted-patient ?d ?p) **(acted-unprofessionally ?d)**))

(:action discuss-work-pressure
:parameters (?d1 ?d2 - doctor ?g - generalLoc)
:precondition
(and **(acted-unprofessionally ?d1)** ... **(happy ?d2)** ...)
:effect
(and **(shown-pressure-of-work ?d1)** ...))

(:action dismiss-work-concerns
:parameters (?d1 ?d2 - doctor ?g - generalLoc)
:precondition
(and **(acted-unprofessionally ?d1)** ... **(angry ?d2)**))
:effect
(and **(shown-pressure-of-work ?d1 ?d2)**))

(:action harass
:parameters (?c1 ?c2 - character ?g - generalLoc)
:precondition (and ...)
:effect (and **(angry ?c2)** (harassed-colleague ?c1 ?c2) ...))

(:action arrange-date
:parameters (?c1 ?c2 - character ?g - generalLoc)
:precondition (and ...)
:effect (and **(happy ?c1)** ...))

Segment subgoal	(shown-pressure-of-work DrAdams DrMiller)
Plan if: (happy DrMiller)	(patient-consultation DrAdams ...) (discuss-work-pressure DrAdams DrMiller ..)
Plan if: (angry DrMiller)	(patient-consultation DrAdams ...) (dismiss-work-concerns DrAdams DrMiller ..)

Figure 3. Narrative modelling example: segment subgoal loosely specified (disjunctive specification [8]) allowing for different plans to be generated depending on state of the narrative world: if Dr Miller is happy then her response to Dr Adam's failure to cope with pressures of work (sympathetic discussion) is very different to if she is angry (dismisses his concerns).

5 Narrative Generation

Our approach to multiplot narrative generation is based on an incremental heuristic search through the space of partially ordered authored subgoals for each of the subplots. Within each loop of the search the following are selected: the next subplot to generate a narrative segment for (ie which subplot to "transition" to), and which subgoal on the frontier of the orders for that subplot to use to generate that narrative segment. Then the narrative is generated forwards

① PDDL3 Ordered Subgoals for SINGLE subplot ②

```
(sometime-before
   (received-treatment Roberts)          ;sg2
   (taken-case DrGreen Roberts))         ;sg1
(sometime-before
   (patient-incident DrGreen Roberts)    ;sg3
   (received-treatment Roberts))         ;sg2
(sometime-before
   (emergency-treatment Roberts)         ;sg4
   (patient-incident DrGreen Roberts))   ;sg3
(sometime
   (discussed-ethics DrGreen DrCook))    ;ex1
(at-end
   (patient-outcome DrGreen Roberts))    ;goal
```

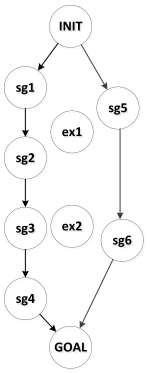

Figure 4. ① PDDL3.0 specification of subgoal orders for a single subplot using the sometime-before and sometime modal operators ② Graphical representation of the partial order showing: all subgoals are ordered before the goal; {sg1, ..., sg4} and {sg5, ..., sg6} are totally ordered (sometime-before modal operator); ex1, ex2 can occur any time (sometime modal operator).

by generating that segment.

One challenge is to achieve even distribution of subplots over the length of the narrative whilst ensuring that adjacent segments are from different subplots (well-formed). Hence the next subplot selected to transition to is the subplot with the most remaining target segments and different to the current.

Another challenge is to select the next subgoal for segment generation that offers the best possibility of matching the input target lengths for the whole narrative. To assess this, the length of the narrative generated so far is considered along with the remaining targets for all subplots and the possibility afforded by each of the frontier subgoals for the selected subplot.

5.1 Multiplot Generation Algorithm

An outline of our multiplot narrative generation algorithm is shown in Fig. 5. The input is (A, I, T, L, G, S) as follows: A, a set of pre- and post-condition planning actions; I, the initial state of the narrative world; T, a set of target segment counts, where $t \in T$ gives the target for each subplot; L, the target length of the output narrative in terms of the total number of segments it contains; G, a set of goal conditions; S, a partial order over subgoals for each subplot.

A narrative plan is a sequence of actions that maps the initial state into a state where all G are true. The aim is that narratives match the target length, L, contain the target number of segments for each subplot, T, and are well-formed (true if adjacent segments are from different subplots).

As shown in Fig. 5, GENERATE-MULTIPLOT builds up narratives incrementally. In each loop the first step is selecting the next subplot to switch to, procedure NEXT-SUBPLOT [line 6 and below], then to select the next subgoal to use for generating the next narrative segment, procedure NEXT-SUBGOAL [line 7 and below]. Once a subgoal has been selected this is used to generate the next segment of the narrative using the current state of the narrative world as the initial state and the selected sg as goal, GENERATE-PLAN [line 8]. This small portion of narrative is generated using the classical planner METRIC-FF [12] which is embedded in the system. At this point the generated segment is visualised to the user [line 9], via staging in a 3D world

1: **procedure** GENERATE-MULTIPLOT(A, I, T, L, G, S)
2: $C \leftarrow I$
3: $Narr, last\text{-}sp, seg, sg, sp \leftarrow \{\}$
4: $\forall s \in S : count_s \leftarrow 0, N \leftarrow 0$
5: **while** $G \neq \{\}$ **do** ▷ Loop
6: $sp \leftarrow$ NEXT-SUBPLOT($S, last\text{-}sp$)
7: $sg \leftarrow$ NEXT-SUBGOAL($S, sp, count_{sp}, T$)
8: $seg \leftarrow$ GENERATE-PLAN(C, A, sg)
9: VISUALIZE-SEGMENT(seg)
10: $C \leftarrow$ADVANCE-STATE(C, seg) ▷ Apply plan actions
11: $G \leftarrow$UPDATE-GOALS(C)
12: $Narr \leftarrow Narr \bullet seg$
13: $count_{sp}$++, T_{sp}--, $last\text{-}sp \leftarrow sp, N$++
14: **end while** ▷ Until all goals solved
15: **end procedure**
16: return $Narr$

Figure 5. Outline Algorithm. Multiple subplot narratives are generated incrementally (while loop line 5). In each loop the next subplot to switch to is selected, and from that the next subgoal to use to structure the next segment of the narrative. As each segment is generated it is visualized to the audience (staged in a 3D world). For more details see text.

(as shown in Fig. 1). Then the current state of the narrative world is advanced by application of the narrative actions in seg, ADVANCE-STATE [line 10] and the overall narrative, *Narr*, is extended forwards by concatenation of seg. The main loop continues until the termination condition: a narrative state with all goal conditions are achieved.

Note that narrative generation operates in a "plan-execute" loop (as shown in Fig 1), with narrative segments generated and then visualized incrementally. Hence backtracking isn't possible but this is managed via loose sub-goal specification (as detailed in section 4). It allows for narrative continuation regardless of narrative state: something of particular importance in interactive systems [22].

Procedure NEXT-SUBPLOT
The subplot sp for which the following are both true:
1. the most segments still remaining to be interleaved into the narrative i.e. the greatest difference between target count, T_{sp}, and the count of segments, $count_{sp}$ already in *Narr*
2. sp is different to the previous loop: $last\text{-}sp \neq sp$.

Procedure NEXT-SUBGOAL
Subgoal, sg, is the one that minimizes D=L-(N+N_h) where, L is the target length for the narrative, N is the number of segments in the narrative generated so far and N_h denotes the number of segments still to be generated for all subplots.

The value N_h is calculated for each unvisited frontier subgoal in the partial order S as the total of the following distances: for subplot sp, the distance of the path from sg to $goal$ in S_{sp}; and for each other subplot, the remaining target distance in T (decremented as segments are added [line 13]).

For subgoal sg, if the distance D doesn't provide the potential to match the target length L, there are two considerations:
1. if $D < (L - (N + N_h))$: any *exposition* subgoals in the partial order S_{sp} are also considered now with the subgoal returned by NEXT-SUBGOAL drawn at random from the set of all exposition subgoals and sg. This enables matching the overall target length by adding additional exposition relevant to the subplot theme.
2. if $D > (L - (N + N_h))$: it is also possible to reduce the subplot length by shortening the path to match the desired length target. However, in practice, for experiments with our prototype this wasn't required (see section 7.1).

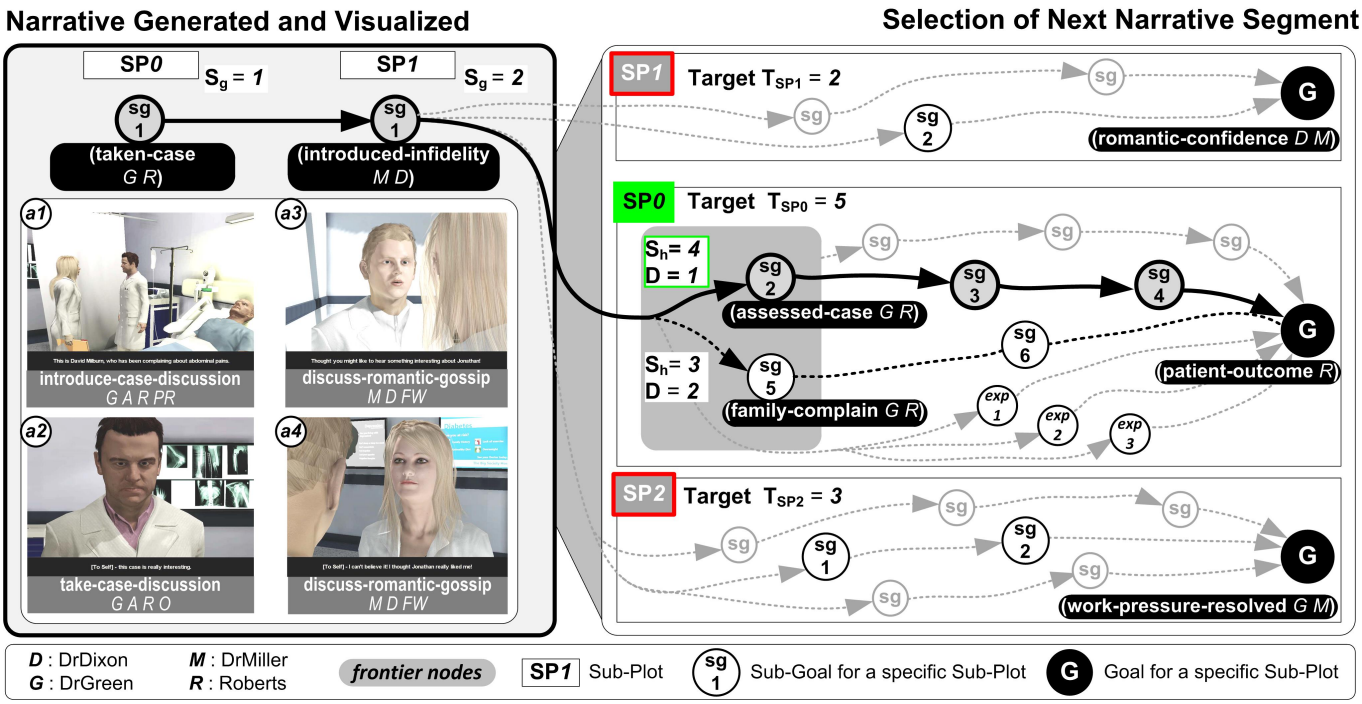

Figure 6. Example of Multiplot Generation.The figure shows the situation after the first two narrative segments have been generated and visualized (the actions labelled A1, A2 for SP0 and A3, A4 for SP1 in the box on the left hand side). At this point the narrative generator must select: 1) which subplot to transition to; and 2) which subgoal to use to generate the next segment. In the figure subplot SP0 (highlighted green) is selected as the subplot to transition to. For selection of the next segment, the frontier narrative subgoals SG2 and SG5 are considered, SG2 provides the closest match to the target narrative length but it doesn't match exactly so the exposition subgoals are also considered with one of these selected (ties broken randomly). Further detail: see section 6.

6 Narrative Generation Example

As illustration consider the example shown in Fig. 6, with three subplots (labelled SP0-SP2), with goals *patient-outcome, work-pressure-resolved* and *romantic-conflict-resolved*, target length, L=12 and initial target subplot counts, T=$\{6, 3, 3\}$. The figure shows the situation after two iterations of the algorithm with the first two segments of the narrative generated using the following subgoals:

(taken-case DrGreen Roberts) ;; SP0-SG1
(introduced-infidelity DrMiller DrDixon) ;; SP1-SG1

and illustrated with thumbnails of the actions visualization (labelled (a1)–(a4)). At this point target counts for the 3 subplots have been adjusted to T to $\{5, 2, 3\}$. Then the subplot for the next narrative segment is selected: in this case SP0 since it has the largest target length (currently $T = \{5, 2, 3\}$) and is different to the previous segment (SP1). The next subgoal is then selected from the frontier of the partial order for this subplot in S. The subgoals considered are:

(assessed-case DrGreen Roberts) ;; SP0-SG2
(family-complain DrGreen Roberts) ;; SP0-SG5

The subgoal that offers the best potential to generate a narrative that meets the target length, L, is then selected based on: the length of the narrative generated so far, N and the combined length of all the segments still to be generated on each of the subplots, N_h. This means considering the length of the path to the goal for each of the frontier narrative nodes and combining that with the target number required on each of the other subplots and then minimizing the distance D=L-(N+N_h). Hence for subplot SP0, for the situation in Fig. 6, where the current targets for SP1 and SP2 are 2 and 3 respectively, L=12 and N = 2 then the values for N_h and D are:

Subgoal	N_h=$Path(sp0)$ + T_{sp1}+T_{sp2}	D=L-(N+N_h)
SP0-SG2	4 + (2 + 3)	12 - (2 + 9) = 1
SP0-SG5	3 + (2 + 3)	12 - (2 + 8) = 2

and since we are minimizing D, subgoal SP0:SG2 is the best. However, at this point, since no values of D equal 0 (i.e. none exactly match the target length), exposition subgoals are also considered. Hence also considered are:

(discuss-ethics DrGreen Roberts) ;; SP0-EX1
(background-patient Roberts) ;; SP0-EX2
(background-doctor DrGreen) ;; SP0-EX3

with the choice being made at random: in this case, selection of subgoal sp0-ex3 as shown in Fig. 6 and that is used as the next goal for incremental narrative generation. This process continues till a state of the narrative world is reached where all the subplot goals are true.

7 Experimental Evaluation

Our prototype interactive narrative was used in the experiments. The narrative domain model for our virtual hospital environment has 15 different locations and 8 doctors, 5 nurses, 4 patients, 4 relatives. For the experiments a test set of narrative planning instances were generated, scaled from 2-4 subplots, with target narrative lengths, L, of 8, 12 and 20 ($\pm$ 2) segments respectively and random assignment of integer target segment counts, T, across subplots. Subplot goals

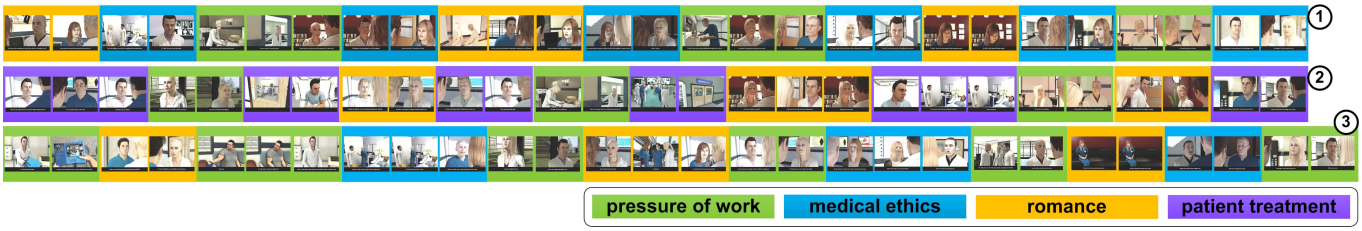

Figure 7. Generated Narrative Thumbnails (narrative targets of 3 subplots, target length 12, target segment counts {6,3,3} The alteration of colours gives a simple overview of the balance and pace between subplot in each different narrative variant. Hence ① has 6 segments on medical ethics subplot and 3 on romance and work pressure; ② 6 on patient treatment; ③ 6 on work pressure. All demonstrate: well formed subplot interleaving; even subplot distribution across duration of narrative. Note how the different subplots of a narrative can interact: for example, for narrative ① the "romance" subplot (orange) can impact on the "clinical" subplot (blue) as follows "acts unprofessionally due to romantic deceptions".

Dr Miller

Dr Adams

Dr McNair

Dr Dixon

Dr Laverick

Dr Cook

Nurse Smith

Subplot	Protagonist	Additional Characters
Romance	**Nurse Smith**	Dr Laverick, Dr Taylor, **Dr Miller**
Pressure of work	**Dr Adams**	Dr Cook, **Nurse Smith**
Medical ethics	Dr McNair	**Dr Miller**, **Dr Adams**, Dr Dixon, Nurse Mills

Figure 8. Illustration of Character "Floating" over Subplots. Here we consider the example narrative ① that was shown as thumbnails in Fig. 7. The characters appearing in the narrative are shown across the top and from the thumbnails it can be seen that the characters Nurse Smith and Dr Adams appear as protagonist in one subplot and as a supporting character in another (for Nurse Smith this is the romance and pressure of work subplots, and for Dr Adams the pressure of work and medical ethics subplots). Further, the character Dr Miller appears as support in two of the subplots.

were randomly selected from the set of facts which were tagged as: work pressure, romance, medical ethics or patient treatment.

7.1 Matching Target Narrative Properties

It is important that generated narratives match the target narrative properties and to assess this we we generated narratives for the set of test narrative problem instances. As an illustration of these narratives, Fig. 7 shows some visualization summaries as thumbnails. In addition, Fig. 8 gives an illustration of the interdependence of the different subplots with the appearance of characters in different subplots and in different roles.

With regard to matching target multiplot length, our approach achieved 100% fit to integer target counts. This performance can be explained in part due to properties of the chains of ordered subgoals which are kept small in order to provide flexibility: short subplots are possible as are longer segments via composition of shorter chains (e.g. just the small fragment shown in Fig. 4 can yield subplots from length 3 segments (path through subgoal SGX) to 10 (combination of paths through subgoals SG1, EX1 and SG5).

All generated narratives featured well-formed interleaved subplots, (i.e. adjacent segments are from different subplots). This property is guaranteed by the requirement of procedure NEXT-SUBPLOT that *last-sp* $\neq$ *sp*, providing that the relative subplot target lengths are within the following bound: the difference between the number of segments in the longest subplot and the combined lengths of all others is ≤ 1. Our problem set was generated for target counts within this bound.

Results also showed that the approach achieves the desired even distribution of subplots over the whole of the narrative. As illustration, below are tabulated the mean separation between segments from the *same* subplot (i.e. count of the intervening segments) for 2, 3 and 4 subplot narratives with target counts as shown.

#Subplots (Targets)	2 (T={4,4})	3 (T={3,3,6})	4 (T={5,5,5,5})
Mean Separation	1:1	3:3:1	3:3:3:3

Finally we note that narrative generation with well-formed subplots retains the generative power of single-plot generation (a strength of planning [26]).

7.2 Narrative Generation Performance

The average times for generation of narrative *segments* for our test problem set were 0.75, 1.4 and 3.9 seconds respectively for the 2, 3 and 4 subplot instances.

The increase in generation time is a consequence of the increase in the overall length of the narrative as additional subplots are required to be interleaved. These timings are acceptable for use in our Interactive Narrative since (see Fig. 1), narrative generation and "execution" (i.e. presentation on a 3D stage) operate in parallel, with narrative generation taking place during the presentation of the previous segment. Given that the average duration of narrative segment presentation is 25 seconds, this allows ample time for generation to be accommodated.

8 User Evaluation

We staged user experiments to evaluate how users would perceive multi-plot narratives compared to the standard single-goal, single-plot narrative in the same domain. Thirty adults participated in the evaluation: compensated for their time with an online retailer voucher worth €28. A consistent protocol was used across assessment of User Preference 8.1 and understanding 8.2: all narrative presentations and questions were delivered via an online questionnaire with order of presentation controlled across subjects.

Figure 9. User preference for Interleaved vs Non-Interleaved narrative presentations: it can be seen that the user group expressed a very clear preference for the interleaved multiplot version when asked "which narrative presentation did you prefer?".

8.1 User Preference

Our aim was to explore whether user perception of system generated interleaved narratives reflected what they're used to seeing in human-authored dramas and took user preference as a proxy for this. To this end we adapted the method of [19] (text to 3D visualization), and asked participants to compare different narratives with the same content presented in different order (and with the same semantics) as follows: (i) subplots interleaved using the approach introduced in the paper; (ii) non-interleaved subplots ordered $\{\{s1_1,...,s1_n\},\{s2_1,...,s2_n\},...,\{sn_1,...,sn_n\}\}$.

Participants were randomly allocated to groups to control between subjects for: interleaved vs non-interleaved; and order of watching videos. Users viewed a total of 4 presentations (2 variations of 2 narratives each with 3 subplots) to avoid cognitive overload by introducing too many characters.

The online questionnaire asked participants to: "Please explain your reasons for preferring the presentation" and users entered free text responses which were judged to relate to subplots if explicitly mentioned (or similar e.g. storylines, storythreads). The results are shown in Fig. 9, and show that the majority, 72%, chose the interleaved presentation. In addition, participants were also asked whether they enjoyed the presentations and given the opportunity to provide free text responses: 80% gave positive responses and of those, 65% volunteered the subplots as the reason. The following give a flavour of the participant responses:

"There were multiple story lines all taking place in parallel."
"Stories are told at once rather than one by one, ... more dramatic"
"The layout .. was more dramatic. As and issue was raised then it moved on and didn't resolve it straight away ..."
"they have a few stories intertwined so the viewer doesnt get bored there is always something new to focus on"
"The storylines where mixed with each other rather than continuing until they were finished so the events felt more realistic."

8.2 Narrative Understanding

We were keen to show that the interleaving of subplots didn't harm story understanding. To demonstrate this we used the QUEST approach where narratives are represented as conceptual graphs that are used to rate the relative quality of comprehension questions [7]. The QUEST evaluation consists of a task for users to complete in order to demonstrate story understanding. It was originally developed for text understanding and has been widely used for IS following [7] by using dynamically generated Q/A pairs. It has been used for the same purpose in [10, 26, 14]). Our intention in using QUEST was to demonstrate that plot interleaving did not impair story understanding.

For the study participants were randomly assigned to groups to watch either interleaved or non-interleaved presentations of two narrative (order of narrative viewing was controlled). Afterwards they

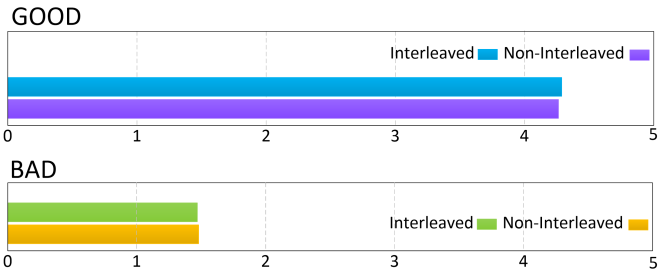

Figure 10. Results of QUEST evaluation: mean responses for GOOD and BAD question-answer (QA) pairs for user groups who watched interleaved or non-interleaved narrative presentations. Understandability is indicated if scores for GOOD QA pairs are high and low for BAD pairs. As anticipated, the results show no difference in story understandability between the different variants (see text for further details).

were asked to assign goodness of answer (GOA) values to question-answer (QA) pairs and then the correlation with the predicted quality given by the QUEST graphs was assessed. For example, a sample pair from our experiments was: Q: *"What did Dr. Thompson do about his patients treatment?"* A: *"He changed the treatment."* and where GOA depends on narrative content i.e. whether treatment changed.

Each user rated the GOA of six QA pairs, selected from the QUEST model, with a value from 1 (very bad) to 5 (very good). Ratings were compared against measures of reachability and arc distance in the QUEST graph for the narrative with expected values for the GOA with 5 (very good) for those with arc distance 1, 4 for those with arc distance 2, and so on, with 1 (very bad) expected for QA pairs that were unreachable in the QUEST graph. After [25] we partitioned the question space into "good" and "bad" QA pairs, where in a rating system of 1–5 good and bad pairs have a system predicted GOA greater or less than neutral respectively and with understandability indicated if scores for bad pairs were low and high for good.

The results were as follows: mean values for good QA pairs of 4.25 and 4.27 out of 5 (for interleaved and non-interleaved respectively); and for bad pairs 1.47 and 1.44. These results indicate understandability and also no difference between interleaved and non-interleaved presentations since in both cases t-test shows no significant difference. This is a promising result supporting our expectation that there is no loss of story understanding when narrative subplots are interleaved.

9 Conclusions

We have extended our earlier landmark approach to plan-based narrative generation to enable it to take into account subplots within the same framework. It should be emphasised that this approach is compatible with previous PDDL based representation of narrative actions and doesn't require bespoke narrative representation.

In the evaluation we demonstrated that the approach is able to generate narratives that conform to different subplot profiles, specified in terms of the number of subplots interleaved and the relative time spent on each presentation. Results of our user evaluation also supported our prediction that generated narratives correspond to user expectations of the genre with the majority of users preferring presentations with interleaved subplots, and with the majority of positive comments attributing this to the subplots.

Overall, these results demonstrate the real potential of the approach to automatically generate interleaved narratives that match user expectations of human-authored narratives and which are clearly preferred by the users.

REFERENCES

[1] Ruth Aylett, João Dias, and Ana Paiva, 'An Affectively Driven Planner for Synthetic Characters', in *Proc. of 16th Int. Conference on Automated Planning and Scheduling (ICAPS)*, (2006).

[2] B-C Bae and R. Michael Young, 'A Use of Flashback and Foreshadowing for SurpriseArousal in Narrative Using a Plan-Based Approach', in *Proc. of 1st Int. Conference on Interactive Digital Storytelling (ICIDS)*, (2008).

[3] David Bordwell, *Narration in the Fiction Film*, Madison: University of Wisconsin, 1985.

[4] David Bordwell and Kristin Thompson, *Film Art: An Introduction*, McGraw-Hill Higher Education (9th Edition), 2010.

[5] Marc Cavazza, David Pizzi, Fred Charles, Thurid Vogt, and Elisabeth André, 'Emotional input for character-based interactive storytelling', in *Proc. of the 8th Int. Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, (2009).

[6] Yun-Gyung Cheong and R Michael Young, 'Suspenser: A Story Generation System for Suspense', *Transactions on Computational Intelligence and Artificial Intelligence in Games*, (2014).

[7] David B. Christian and R. Michael Young, 'Comparing Cognitive and Computational Models of Narrative Structure', in *Proc. of the 19th National Conference on AI (AAAI)*, (2004).

[8] B. C. Gazen and C. Knoblock, 'Combining the Expressiveness of UCPOP with the Efficiency of Graphplan', in *Proc. of the European Conf. on Planning (ECP)*, (1997).

[9] Stephen Gilroy, Julie Porteous, Fred Charles, Marc Cavazza, Eyal Soreq, Gal Raz, Limor Ikar, Ayelet Or-Borichov, Udi Ben-Arie, Ilana Klovatch, and Talma Hendler, 'A Brain-Computer Interface to a Plan-based Narrative', in *Proc. of the 23rd Int. Joint Conf. on Artificial Intelligence (IJCAI-13)*, (2013).

[10] Stephen W. Gilroy, Julie Porteous, Fred Charles, and Marc Cavazza, 'Exploring Passive User Interaction for Adaptive Narratives', in *Proc. of the 17th Int. Conf. on Intelligent User Interfaces (IUI-12)*, (2012).

[11] P. Haslum, 'Narrative Planning: Compilations to Classical Planning', *Journal of AI Research*, **44**, 383–395, (2012).

[12] J. Hoffmann and B. Nebel, 'The FF Planning System: Fast Plan Generation through Heuristic Search', *Journal of AI Research*, **14**, 253–302, (2001).

[13] Tomas Hromada, Martin Černý, Michal Bída, and Cyril Brom, 'Generating Side Quests from Building Blocks', in *Proc. of the 8th Int. Conf. on Interactive Digital Storytelling (ICIDS)*, (2015).

[14] Arnav Jhala and R Michael Young, 'Cinematic visual discourse: Representation, generation, and evaluation', *Computational Intelligence and AI in Games, IEEE Transactions on*, **2**(2), 69–81, (2010).

[15] Michael Lebowitz, 'Story-telling as planning and learning', *Poetics*, **14**(6), 483–502, (1985).

[16] Boyang Li and Mark O. Riedl, 'An Offline Planning Approach to Game Plotline Adaptation', in *Proc. of the 6th National Conference on AI and Interactive Digital Entertainment (AIIDE)*, (2010).

[17] Bradford Mott, Sunyoung Lee, and James Lester, 'Probabilistic goal recognition in interactive narrative environments', in *Proc. of the 21st National Conference on Artificial Intelligence (AAAI)*, (2006).

[18] Mark J. Nelson and Michael Mateas, 'Search-Based Drama Management in the Interactive Fiction Anchorhead', in *Proc. of the st AI and Interactive Digital Entertainment Conference (AIIDE)*, (2005).

[19] Brian O'Neill and Mark Riedl, 'Dramatis: A Computational Model of Suspense', in *Proc. of the 28th National Conference on AI (AAAI)*, (2014).

[20] J. Porteous, F. Charles, and M. Cavazza, 'NetworkING: using Character Relationships for Interactive Narrative Generation', in *Proc. of 12th Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, (2013).

[21] J. Porteous, J. Teutenberg, D. Pizzi, and M. Cavazza, 'Visual Programming of Plan Dynamics using Constraints and Landmarks', in *Proc. of the 21st Int. Conf. on Automated Planning and Scheduling (ICAPS)*, (2011).

[22] Julie Porteous, Marc Cavazza, and Fred Charles, 'Applying Planning to Interactive Storytelling: Narrative Control using State Constraints', *ACM Transactions on Intelligent Systems and Technology (ACM TIST)*, **1**(2), 1–21, (2010).

[23] M. Riedl, 'Incorporating Authorial Intent into Generative Narrative Systems', in *Proc. of the AAAI Spring Symposium: Intelligent Narrative Technologies II*, (2009).

[24] M. Riedl, CJ Saretto, and R.M. Young, 'Managing interaction between users and agents in a multi-agent storytelling environment', in *Proc. of the 2nd Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*. IFAAMAS, (2003).

[25] M. Riedl and R.M. Young, 'Open-world planning for story generation', in *Proc. of the Int. Joint Conference on AI (IJCAI))*, (2005).

[26] Mark O. Riedl and R. Michael. Young, 'Narrative Planning: Balancing Plot and Character', *Journal of AI Research*, **39**, 217–267, (2010).

[27] ER, 1994. (TV Series) Warner Brothers, 1994-2009.

[28] HOUSE, 2004. (TV series, Universal, 2004-2012.

[29] MASH, 1972. (TV series) 20th Century Fox, 1972-1983.

[30] SCRUBS, 2001. (TV Series), ABC Studios, 2001-2010.

[31] Seinfeld, 1989. (TV series) West-Shapiro Productions and Castle Rock Entertainment, 1989-1998.

[32] Manu Sharma, Santiago Ontañón, Manish Mehta, and Ashwin Ram, 'Drama Management and Player Modeling for Interactive Fiction Games', *Computational Intelligence*, **26**(2), (2010).

[33] Greg M. Smith, 'Plotting a TV Show About Nothing: Patterns of Narration in Seinfeld', in *Creative Screenwriting*, volume 2(1), (1995).

[34] T. J. Smith, D. T. Levin, and J. Cutting, 'A Window on Reality: Perceiving Edited Moving Images', *Current Directions in Psychological Science*, **21**, (2012).

[35] Kristin Thompson, *Storytelling in Film and Television*, Harvard University Press, 2003.

[36] Peter Weyhrauch, *Guiding interactive drama*, 1997.

[37] R. M. Young, 'Creating Interactive Narrative Structures: The Potential for AI Approaches', in *AAAI Spring Symposium in Artificial Intelligence and Entertainment*. AAAI Press, (2000).

[38] R. Michael Young, 'Notes on the use of plan structures in the creation of interactive plot', in *Proc. of the AAAI Fall Symposium on Narrative Intelligence*, (1999).