

Planning Technologies for Interactive Storytelling

Julie Porteous

the date of receipt and acceptance should be inserted later

Abstract Since AI Planning was first proposed for the task of narrative generation in Interactive Storytelling (IS) it has emerged as the dominant approach in this field. This paper traces the use of planning technologies in this area, considers the core issues involved in the application of planning technologies in IS and identifies some of the remaining challenges.

1 Introduction

The goal of Interactive Storytelling (IS) research is the development of multimedia systems in which users can interact and influence, in real-time, the evolution of a narrative as it is presented to them.

AI planning was first proposed for the task of narrative generation in IS by Young (1999) because it provided a natural fit for representation with narratives as plans; ensured causality which is important for the generation of meaningful and comprehensible narratives; and provided considerable flexibility and potential generative power. Hence it is not surprising that during the intervening time the plan-based approach has been enthusiastically adopted (for example, the systems reported by Cavazza et al. (2002a), Aylett et al. (2006), Riedl and Young (2010) and Porteous et al. (2013)).

In plan-based IS systems a plan generation engine is embedded within a planning and execution loop where the planner generates an initial narrative which is “executed” via presentation to the user (e.g. using 2D or 3D graphics [Mateas and Stern (2005), Porteous et al. (2013)], filmic content [Piacenza et al. (2011)] or text [Orkin and Roy (2012)]) and where user interaction can change the state of the planning world and hence require replanning or repair.

Julie Porteous
School of Computing,
Teesside University, U.K.
E-mail: j.porteous@tees.ac.uk

Given the real-time nature of IS systems, planner performance is clearly an important issue, however our intuition is that this is less of an issue than in classical planning benchmark domains (as illustration, contrast the 2 million nodes explored in the narrative example given in [Riedl and Young (2010)], with the astronomically large search spaces solved by heuristic search planners – 10^{100} is not unheard of [Hoffmann (2011)]) and that other aspects, such as narrative control and authoring, are more central.

These core issues of planning and narrative are reflected in the content of this chapter: we start in the next section with an overview of planning and IS systems to date, discuss issues to do with Narrative Planning Models (ie domain and problems) and identify a minimal set of representational assumptions which underpin all applications of planning to IS (section 3) and for the remainder of the paper the focus is two enduring challenges for planning in IS, namely Narrative Control (section 4) and Authoring (section 5), which are discussed and state-of-the art solutions reviewed.

2 Plan-based Narrative Generation

Young (1999) first proposed the use of AI planning for generation of narrative in IS since it offers a rich representation of the causal structure of output narratives along with the story diversity that results from its generative power. Since then plan-based approaches have been enthusiastically adopted.

Early work in IS tended to feature older planning approaches, such as partial-order planning [Riedl and Stern (2006); Young (2000)] and Hierarchical Task Networks (HTN) [Aylett et al. (2006); Cavazza et al. (2002b); Hoang et al. (2005); Karlsson et al. (2006)]. Indeed these approaches still feature in some current systems. For example, the LOGTELL system features a partial-order planner for initial plot-generation [Da Silva et al. (2010)], the intent handling in the IPOCL system is embedded in a partial-order approach [Bae et al. (2011)] and HTN's feature in the work of [Kelly et al. (2007); Paul et al. (2011)].

The continued use of partial-order planning in IS can to some extent be attributed to the flexibility that it offers. For example, it was shown to provide leverage with respect to aspects such as character intent and believability [Riedl and Stern (2006); Riedl (2004)]. However, this must be weighed against the poor performance of partial order planners (especially when compared to state-of-the-art state-space planners). As an illustration of this, experiments on narrative generation with the IPOCL planner report a time of approximately 12.3 hours to generate an example narrative plan [Riedl and Young 2010]. It is not clear that such an approach could scale to a realistically sized system or be able to perform within the constraints of a real-time interactive system.

The enduring popularity of HTN's can be explained to some extent by the fact that IS has been viewed as a case of knowledge-intensive planning, whereby a vast amount of prior knowledge needed to be engineered into the planning domain, and this justified the use of specific formalisms such as HTN

planning. However, control knowledge can be difficult to specify and maintain precisely because it is embedded [Cavazza et al. (2002a)].

More promise is offered in the recent trend in IS to the use of forward state-space planning approaches in the spirit of HSP [Bonet and Geffner (2001)] and FF [Hoffmann and Nebel (2001)] for use in IS: for example, Barros and Musse (2007) used Metric-FF Hoffmann (2003), Pizzi et al. (2007b) used a variant of HSP in combination with the use of RTA* [Korf (1990)], and Porteous et al. (2011a) featured temporal planning with CRIKEY [Coles et al. (2009)]. The explanation for the adoption of these approaches stems mainly from their spectacular performance gains (e.g. as demonstrated in the bi-annual ICAPS International Planning Competition), along with convergence on an expressive representation language, PDDL [Fox and Long (2003)].

Independently of the particular planning algorithms adopted, there has been a tendency to specialise the planner on one particular aspect of IS. In the main, planning has been applied to the task of story generation, for example [Porteous et al. (2011a); Young (2001)]. However, planning has also been applied to the task of control, in IDA [Magerko (2007b)] and to the planning of discourse aspects [Jhala and Young (2005); Young (2007)].

3 Plan-based Representation of Narrative Worlds

This section discusses some key issues in plan-based representation of narrative worlds and identifies a minimal set of key representational assumptions that underpin all applications of planning to IS.

3.1 Duality of Character and Plot

A narrative can be considered from the perspective of its featured characters or from a more abstract perspective at the global level of the plot. In IS the tendency has been to distinguish between character- and plot-based approaches and for planning this has different consequences for representation as discussed below. However a plan-based approach is neutral with respect to this duality since both stances are compatible with the approach at the representational level which can be seen as a spectrum: at one end the focus is on defining autonomous character behaviour for planning from which story will emerge; and at the other end are approaches that use plot requirements to specify the necessary plan based character behaviour.

- *Plot-based Approaches*

With a plot-based approach the representation of the narrative world is based on a model of the baseline plot itself. As a consequence, narrative actions are modelled from a plot perspective and their “execution” can involve one or multiple characters. This is a centralised approach that offers a mechanism for narrative control and hence provides a better solution to the problem of

authorial control in order for improved narrative coherence. However, the control provided by a plot-based approach comes at a price: the potential for individual actions to become over-complicated, along with a reduction in the generative power that more distributed approaches might provide.

- *Character-based Approaches*

The alternative perspective is one which is strongly centred on characters and hence referred to as character-based [Cavazza et al. (2002b)]. The key advantage of this approach stems from the promise of greater flexibility and potential for generative scalability when narratives emerge from the interactions between different narrative characters and their roles. However, these advantages must be traded-off against the loss of narrative control inherent in a character-based approach which can make it difficult to ensure that generated narratives display desired criteria, such as aesthetic properties.

3.2 Narrative Plans

For a large number of story genres, ranging from classic folk tales to modern day stories of heists and bank robberies, the structure of the plot itself can be quite naturally modelled using a plan. In such cases, the goals of the narrative and the actions that characters can take to achieve them are isomorphic to the goals and actions of the plan. In IS, instances of this type of strong connection between the narrative structure and representation have played an important role: from the “bank robbery” scenario discussed by Riedl et al. (2003) through to current systems such as the “riddle” genre of Barros and Musse (2008).

However, it is also possible to make weaker representational assumptions, without implying such a strong connection between the narrative structure and the plan-based representation. For example, in the Madame Bovary inspired system of Pizzi et al. (2007a) a planning representation was used to support the central characters role rather than problem-solving activity based on primitive intentions and desires and then planning was used to enforce causality in sequences of narrative actions intertwined with user interactions.

3.3 IS Problems from a Planning Perspective

- *Plan Quality Criteria*

The question as to what constitutes criteria for “good” narratives in IS is still open. Current state-of-the-art systems require input from human authors such as assigning interesting authorial goals [Riedl (2009)] the design of appropriate initial and goal states [Riedl and Young (2005)], or describing desired tension arcs [Porteous et al. (2011b)]. Minimum length criteria can also play a part, where this implies a tight narrative without superfluous actions. In addition to a quality measure, narrative plans must be produced with hard constraints

on their plausibility: every action seen to be performed by a character must appear to have some motivation or intent [Riedl and Young (2010)].

IS domains differ markedly from the sorts of benchmarks that have featured in planning research, such as the benchmark domains used in the ICAPS International Planning Competitions (ICAPS). Indeed, there are some key aspects of these IS domains that go against the accepted planning wisdom: in IS domains optimality (e.g. in terms of makespan) is not essential, rather it is the “shape” of the plan trajectory that is of central importance; goals are not necessarily “real” goals in the classical sense; and in certain situations the system can actively require plan failure and backtracking. Consider the criteria used to assess plan quality, which sets IS domains apart. Historically, optimality in terms of plan length was the criterion for assessing the quality of plans [Chapman (1987)] with more recent extensions to enable the specification of plan metrics [Hoffmann (2003)] and adequacy criteria such as preferences and constraints [Baier et al. (2007); Chen et al. (2006)]. In IS, quality criteria differ since re-planning is frequently required in response to user interaction, interesting narrative events may be missed if the emphasis is on generating optimal solutions, and there is frequently a need to generate sub-optimal trajectories in order to display narrative concepts such as suspense. Indeed, there is growing consensus that for these types of domains quality criteria should be more concerned with the dynamics of the plan, the shape of its trajectory, in other words, the intermediate states that will be traversed when it is executed.

An example of tool support for the authoring of narrative plan quality criteria via a visual programming interface is discussed in section 5.3.

- *Narrative Planning Goals*

For the application of planning to IS, the nature of goals differs from classical domains since goals need not equate to the end of the narrative (i.e., serve as the narrative driver) and goals can evolve over time. For instance, the GADIN system of Barber and Kudenko (2007) featured dynamic selection and re-selection of story goals in response to narrative dilemmas, the continual multi-agent narrative planning approach of Brenner (2010) allowed for temporary sub-goals, which could be violated in the final state of the narrative plan and included a notion of dynamic sub-goal activation and the Brain-Computer interfaced system reported by Gilroy et al. (2013) enabled the narrative goal to evolve over time depending on users’ affective responses.

- *Role of Plan Failure*

The role of plan failure in IS is very different to that in typical classical planning domains: indeed, in certain situations, an IS system might encourage failure in order to achieve some dramatic effect. In general, planned narratives can fail as a result of user actions. Young (2007) distinguished between actions which don’t interfere with the planned sequence and exceptional user actions which require some form of intervention to avoid failure. The model for recovery, proposed by Riedl and Stern (2006) and Riedl et al. (2003) featured tiered re-planning with mixed off-line analysis of potential sources of inconsistency,

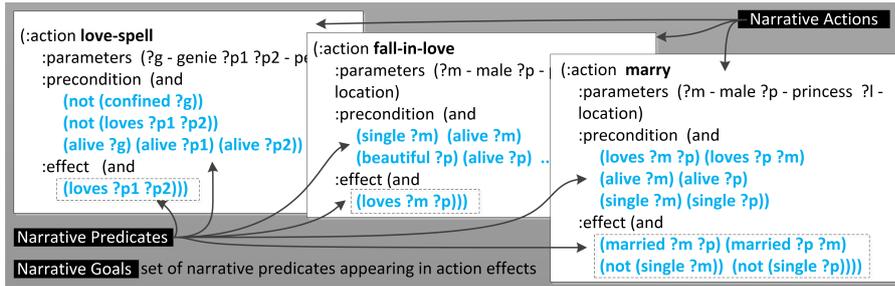


Fig. 1 Illustrative operators from our Aladdin domain represented using PDDL. Operators are specified in terms of pre-conditions (conditions that must hold for an operator to be visualised in the 3D environment) and post-conditions (effects of the operator). Operators feature typed parameters, such as ?g, ?p1, etc, which are bound to constants at run-time.

on-line repair that attempts to restore causal coherence and if all else fails, the system attempts to recover by re-planning.

Synergistically, forward state-space search results in benefits for fast re-planning since the complete state of the world is always known thus facilitating re-planning directly from the point of failure. This in part provided motivation for the decomposition approach to narrative generation introduced in [Porteous et al. (2010)], where narratives are generated via forward search through a series of smaller sub-narratives which are produced incrementally, thus reducing wasted effort should re-planning be required (see section 4.2).

3.4 Minimal Representational Assumptions

Based on the previous discussions a minimal set of representational assumptions can be identified for the development of narrative domain models for planning in IS. These consist of:

Narrative Predicates: a set of predicates which can be used to describe states of the narrative world. These can correspond to object properties, location or properties of characters, or characters’ mental states, including emotional ones.

Narrative Actions: which can be used as planning operators for the generation of narrative plans. These specify the ways in which the state of the narrative world can be changed and are expressed as pre- and post-condition operators.

Narrative Goal: a goal condition, expressed as a set of narrative predicates and which represents the default end state for a scene or for the narrative at large (this can be a default goal to be used in the absence of user interactions which might necessitate it being changed).

As an illustration, some examples of these components for a narrative planning domain based on the Aladdin folk tale are shown in Fig. 1. In the absence of any standardised benchmarking narrative planning domains this is used as

illustration due to the familiarity of the tale, its previous use in narrative research (e.g. [Riedl and Young (2010); Teutenberg and Porteous (2013)]).

4 Control in Plan-based Interactive Storytelling

4.1 Intent-based Narrative Planning

An important factor in narrative understanding is that virtual characters must be believable – the audience must suspend belief and the actions of virtual characters must not threaten this [Bates (1994)]. In other words the audience must perceive them as intentional agents [Dennett (1989)]. Multiagent simulation based approaches to narrative generation neatly solves the problem of endowing agents with intentionality: by treating each character as an autonomous agent with its own beliefs and intents, believable character interactions can emerge from a simulation [Aylett et al. (2006)]. The drawback for such systems is in the fulfilling of global narrative goals and author preferences since a character-only simulation relies on emergence and there is no guarantee of even an approximation of criteria other than believability.

In contrast, plan-based narrative generation approaches can struggle with producing an appearance of intentionality. In particular, plot-based narrative generation by a single classical planner that has available to it all character actions and plot events is able to compose sequences of character actions to fulfil the global goals that make no sense from any individual character’s perspective. To make planning suitable for IS, intent-based planning has been proposed with the IPOCL planner [Riedl and Young (2010)]. It was originally devised as an extension to a causal link planner that searches the space of partial plans, working backward from the narrative goals. The appearance of intent was ensured by including, in parallel, the planning of frames of commitment that span contiguous sub-sequences of actions within a plan. The intent for a frame is defined by a character and a target fact that they intend to achieve. The backward chaining nature of the search ensures that characters’ actions are causally linked to, and therefore appear relevant to, the final action that closes their frame of commitment. This leads to the generation of narratives where characters’ actions have intent, as though planned and performed by autonomous agents, by explicitly representing intent in the narrative and requiring that each action is assigned to some intent. However, intentionality comes at a price and a notable drawback of the IPOCL planner is the high complexity of the search and long run-time making it unsuitable for real-time IS use. This is further compounded by the fact that IPOCL has a non-standard representation language which prevents experimentation with more state-of-the art planning approaches whose performance might be faster.

Interestingly, Haslum (2012) has recently introduced a compilation for IPOCL narrative domains to classical planning. This meant that the intent-based planning problems could now be tackled by a wide range of classical planners, and for IS it was hoped that this would enable real-time narrative

generation. However, the experiments we reported in [Teutenberg and Porteous (2013)] suggested that narrative generation with the compiled domain models was still too slow to realistically be used in an IS system. Further, the coupling of planning for characters’ goals with planning for narrative goals prohibits the generation of certain narratives, most notably those narratives in which one or more characters become unable to complete their intent.

The problem was to develop an approach that would be capable of performing an extensive search of the space of narratives, in order to be able to generate these types of interesting narratives but which would also be capable of performing within the time constraints of a real-time IS system. The solution, which we articulated in [Teutenberg and Porteous (2013)], was to delegate the task of reasoning about intentions to narrative agents who are responsible for the checking of “narrative action relevance” with regard to their individual intentions. Then a single narrative planner was used to generate narratives making use of agents’ relevant actions. These ideas are fully implemented in a planner, called IMPRACTical (Intentional Multi-agent Planning with

4.2 Narrative Control via Trajectory properties

An enduring challenge in IS is how to control the “shape” of narratives that are automatically generated and in our work we have tackled this via an approach that exploits a meta-level abstraction with respect to both time and causality [Porteous et al. (2010)]. This meta-level is represented using *constraints* – key narrative situations for a domain – that are used as intermediate goals which guide operator selection and that the corresponding trajectory exhibits desired properties (for example, in narrative these can be associated with pace, suspense and so on). Please note that while we use the nomenclature of [Porteous et al. (2010)] we observe the similarity to author goals [Riedl (2009)] and also the similarity to the notion of landmarks [Hoffmann et al. (2004)]. However, unlike landmarks, constraints don’t always have to be made true in order to solve the goal, only to conform to desired plan dynamics.

Our planner is a state-space forward search heuristic planner which uses a partially ordered set of constraints to decompose the process of narrative generation into a sequence of sub-problems, where each sub-problem has a constraint as its goal and the planner generates a narrative for each decomposed sub-problem in turn. If the planner can’t generate a narrative for a sub-problem then it simply continues with the next constraint, thus ensuring planner continuation. Once all sub-problems have been tackled, a final narrative can be assembled by composition of the narratives for each sub-problem. However, when our planner is integrated within an IS system a complete plan need not be output in the traditional sense. Instead operators are sent one at a time for 3D visualisation to a user.

In our implemented systems (e.g. Merchant of Venice [Porteous et al. (2010)] and NETWORKING [Porteous et al. (2013)]) this narrative planner operates within a control loop that handles constraint and problem instance

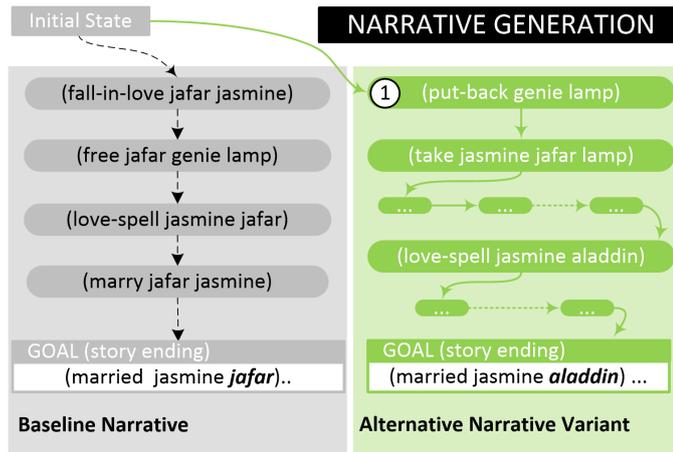


Fig. 2 Example of Narrative Generation. The figure shows the process of narrative generation with an initial baseline narrative (down the left hand side) generated using a domain modelled around a baseline story. Also shown in the figure is the generation of an alternative narrative variant (down the right hand side) which demonstrates the need for additional narrative actions (here, this includes actions such as putting the genie back in the lamp, labelled ①), for a very different story ending with princess Jasmine married to Aladdin).

selection at run-time. Constraints for a particular narrative world are represented using PDDL3.0 and form a partially ordered set of predicates. For a particular planning instance a subset of the constraints are selected (e.g. based on user preferences and interaction history). The planner then uses this to drive narrative generation. Initial problem instances are created for each session on the basis of factors such as user preference and enforced variation between user sessions.

5 Authoring of interactive Narrative Planning Models

IS systems are dynamic environments within which virtual agents act under the control of system generated storylines but where real-time interference with the ongoing story can impact on subsequent narrative unfolding. AI planning is well suited for the task of generating stories in such systems however the creation of the domain models for these planning based story generators raises an important practical domain modelling problem: how to author them to ensure they contain sufficient narrative content in order to produce alternative narratives to the baseline plot. Such alternatives can be part of either narrative generation (i.e. the production of multiple story variants) or interactive narrative in its strict sense (i.e. the system responding to dynamic changes made to the story world).

To illustrate this consider the examples generated using our model of the Aladdin folk tale as shown in Fig. 2 and 3. An initial narrative generated with a domain modelled around a baseline plot might proceed as shown down

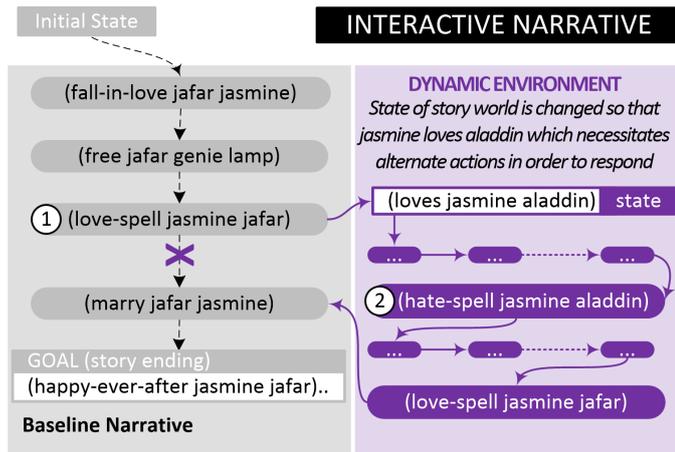


Fig. 3 Example of Interactive Narrative. The figure illustrates the need for alternate actions in an interactive narrative in order to respond to dynamic changes made to the story world: here dynamic changes mean that the love spell (labelled ① which should result in princess Jasmine falling in love with Jafar) goes wrong (with the result that Jasmine loves Aladdin) and consequently alternate actions are required to respond appropriately from that point onwards (such as casting a spell so that Jasmine hates Aladdin ②).

the left hand side of Fig. 2 with characters falling in love, casting love spells and marrying. However, for the system to produce multiple story variants, additional actions are needed for storylines that deviate from the baseline (right hand side of Fig 3). In addition, to be able to respond to dynamic changes to the story world and to continue the story through to the intended ending, alternative actions are required to generate an alternative course of action from that point onwards (Fig. 3).

5.1 Manual Domain Authoring

To date the authoring of such interactive narrative domain models has been largely handled manually, a common strategy being to build up the model via systematic consideration of alternatives around a baseline plot. We termed this approach “narrative variant IS” and have used it in our work for example with *The Merchant of Venice* [Porteous et al. (2010)].

This method explicitly attempts to create interactive variants of pre-existing stories, often classics, for which scholarly analyses exist that facilitate formalisation. This strategy of modifying existing narratives can be reduced to a form of remediation [Bolter and Grusin (1999)], where the development of a narrative domain starts by modelling a default story as it is traditionally described i.e. the baseline plot. This default story is represented as a plan, whose goal state corresponds to the default story ending and each main narrative action is represented as a planning operator. There is no assumption that the story itself is equated to a plan for achieving a specific goal. Even so, describing the

goal state at different levels of detail ensures that various endings are still possible for the same baseline plot. In this context, there exist two determinants of variability around the baseline narrative:

- Run-time modification of the planning domain through user interaction (e.g. removal of facts will invalidate pre-conditions which would otherwise be satisfied, forcing the story along different avenues)
- Off-line modification of the domain model with addition of new narrative actions that are different but compatible in spirit with the baseline plot.

Without such actions, removal of certain facts through interaction may simply result in a dead end, meaning that interaction naturally requires a richer set of actions. Furthermore, even for a given world state, the availability of several compatible actions creates potential for variability triggered by initial conditions or a number of mechanisms for dynamic action selection, such as arbitrarily breaking ties between similar heuristic values in the planning process. The compatibility of operator variants stays within the same semantic context, but this does not mean that the individual operators are directly interchangeable, rather that the generation process will have the possibility to generate varied sequences of operators. Representing narrative actions in the form of planning operators makes explicit their agent-based structure, identifying actors as well as key resources for action – this constitutes a good starting point for the identification of domain predicates which form part of alternative, possibly competing, actions as well as potential targets for interaction.

An indication of the number of different possible narrative variants can be obtained by estimating the size of the narrative space. This depends on the expected narrative path length, n , and the average branching factor, b , (i.e., the number of applicable narrative actions) and can be expressed as b^n . In turn, the branching factor depends on the number of characters in the narrative, c , and the average number of actions that each character has available to them, a , giving, $b = c * a$. As an example, the domain for our Merchant of Venice system featured 3 central characters with average narrative length of 23 actions and an average of 900 instantiated actions for sub-plots selected from the baseline play [Porteous et al. (2010)]. Novel variants around the baseline could be generated in a number of ways, for instance the operator set contained actions which enabled the plot to unfold in different ways to the baseline (e.g., Shylock could choose to show mercy in the final trial, and in so doing avoid the original tragic story ending) and also actions that could be re-combined in novel ways, resulting in new semantics.

5.2 Automated Domain Authoring

Planning Domain Model creation is challenging, time consuming and an obstacle to the further fielded application of planning technology [Zhuo and Kambhampati (2013)]. Hence its no surprise that automation of the process is a topic of current interest both in the AI planning and Narrative research communi-

ties. With respect to AI planning, recent work has been aimed at learning planner action models from correctly observed plan traces (e.g. [Amir and Chang (2008); Cresswell and Gregory (2011); Zhuo et al. (2010)]). However, this work has limited application for narrative domains which do not share the same consistency and alignment with real-world domains as do more traditional planning domains such as “logistics” or “rovers”.

Similarly, in narrative research there has been increasing work aimed at automated creation of content. A popular approach has been the gathering of story elements via crowdsourcing, an approach which can yield abundant content. For example, with the SCHEHERAZADE system Li et al. (2013) employ this approach to acquire typical story elements which can be assembled as plot graphs and used in a process of story generation. With SCENARIOGEN [Sina et al. (2014)] crowdsourcing was used to gather a database of scenarios of everyday activities and likely replacements for use within a serious game context. Orkin and Roy (2012) used a crowdsourcing approach for the hand annotation of logs of the *Restaurant Game* for subsequent use in automating character interactions with human participants in a dialogue-based narrative setting.

An alternative to crowdsourcing aims to obtain narrative content through mining of weblogs and story corpora: the SAYANYTHING system of Swanson and Gordon (2012) selects narrative content on-the-fly from a corpora of weblogs in response to user text-based interaction; whilst the approach of McIntyre and Lapata (2009) attempts to generate narratives using knowledge mined from story corpora for a particular genre.

In our work we were motivated to explore automation of our narrative variant approach and have developed an approach to the automated generation of “missing” actions in the domain model that represent opposites of the baseline ones. We developed an automated approach to the identification and creation of such actions, one which can be seen as automating a manual strategy of considering opposite actions and predicates, thus enabling the automated creation of alternate narrative content that departs from a linear baseline narrative [Porteous et al. (2015)]. Our implemented prototype system, named ANTON, takes as input a planning domain model, identifies missing contrary elements, and uses this to generate new content which is output in an extended version of the domain model. Since it is important that the ANTON generated content is human-readable (e.g. to trace stories during development) and given that contrary relationships form the basis of our approach to action generation, automatic labelling of new content in ANTON is based on antonyms, sourced from a range of linguistic resources such as [WordNet] and [Merriam-Webster].

5.3 Authoring Tools: Visual Programming of Trajectory Properties

Earlier we reflected on the nature of quality criteria for narrative plans and observed that these are more concerned with the dynamics of plan trajectories – the intermediate states that will be traversed when it is executed – than the

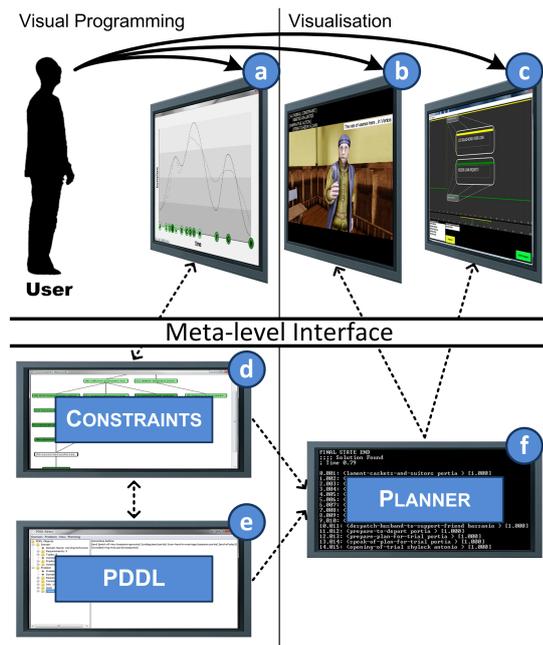


Fig. 4 Visual programming of plan dynamics: user interacts via narrative arc window, drawing desired arcs, adjusting tension levels (a); global properties of generated narratives assessed via visualisation windows (b) and (c); lower level components accessed via interface (d) and (e); user invocation of planner using specified plan dynamics (f)

more traditional classical planning domains. For these trajectories different application-specific interpretations apply: for example in new media they relate to narrative features such as suspense, pace and so on.

Based on these observations, in our work we addressed the problem of how to provide authoring support for the specification of the quality criteria for narrative domains. Since controlling trajectory by direct manipulation of the domain model is a considerable challenge and because trajectory suits a visual representation our solution was to use a visual programming approach. Building on the meta-level of representation of our narrative planning approach (discussed in section 4.2) we developed a visual representation of this in the form of a narrative arc for the authoring and exploration of plan dynamics (as reported in [Porteous et al. (2011b)]).

An overview of our developed system architecture is shown in figure 4. Users interact at the meta-level via a Narrative Arc Window (a) and can also explore generated narratives via two visualisation windows: an Animation Window (b); and a Timeline Window (c). The system also features some hierarchically organised lower level components. They include the constraints (d) and other PDDL constituents of the domain model (e). Invocation of the planner (f) is driven by the user and enables them to explore the narrative possibilities of different sets of plan dynamics.

6 Conclusions

Despite the wide-scale adoption of planning, other techniques have been proposed for narrative in IS such as Branching Narrative Structures [Hill et al. (2001); Spierling et al. (2006)], Behaviour Trees [Champandard (2008); LLansó et al. (2009)], Bayesian Networks [Arinbjarnar and Kudenko (2010); Sparacino (2003)], Case-based Reasoning [Gervás et al. (2005); Swanson and Gordon (2010)], Assumption Based Truth Maintenance Systems [Sgouros et al. (1996)] and Monte-Carlo Tree Search [Kartal et al. (2014)].

Nevertheless, on the basis of the following key aspects, our conclusion is that narrative generation for IS is best supported by AI planning:

- *Causality*

The maintenance of local and global story causality is an important dimension of narrativity, as argued by Ryan (2006), which has been observed to be an important factor in user experience [Trabasso and van den Broek (1985)]. Hence causality was part of the original rationale for proposing a planning-based approach to narrative in IS [Young (2000)] – the ability of the technology to provide appropriate computational narrative structures, cognitive models and support for shared user-system control of narrative development. The mechanism used in planning algorithms to enforce causal consistency is through the achievement of required pre-conditions.

- *Generative Power*

Regarding shared control of an interactive system, as highlighted by Young (1999), another strength of planning is its generative power – as a consequence of which, it is able to re-plan in a timely fashion when user actions necessitate changes to the planned narrative. The generative power of a plan-based approach also provides a powerful mechanism to scale up to realistically-sized systems [Charles and Cavazza (2004)].

- *Story Progression*

In IS, there is a requirement for story progression in the absence of user interaction. This is in contrast to the situation with games that feature narrative content, such as HITMAN [IO Interactive (2000)] and HEAVY RAIN [Cage (2010)], where the driver for progression comes solely from user interaction, in other words with no user interaction there would be no progression. A key strength of planning, unlike other technologies, is that it provides a driver for story generation which allows a story to progress regardless of user interaction.

- *Real-time Response*

Some researchers have proposed the use of reactive planning languages such as ABL [Mateas and Stern (2002)] for content generation “on-the-fly” and have suggested that classical planners lack direct support for this during user interaction via different interaction modalities [Magerko (2007a)]. However, as already argued, this is not the case for state-of-the-art forward state-space planners, such as HSP [Bonet and Geffner (2001)] and FF [Hoffmann and Nebel

(2001)] which build a plan forwards: an approach that actively supports narrative generation in the presence of user interaction. Interestingly, reactive planning approaches have also been advocated in support of interaction based on dialogue, as in FAÇADE [Mateas and Stern (2002)] which is largely based on dialogue, even to the point of subordinating plot evolution to dialogic actions. However the same arguments in favour of planning apply here: the flexible, generative nature of planning makes it possible to respond in real-time to user dialogue without sacrificing global narrative properties [Cavazza et al. (2009)].

- *Flexibility*

An alternative to plan-based narrative generation that has featured in some IS systems, is the use of branching graph narrative structures [Hill et al. (2001); Spierling et al. (2006); Tully and Turner (2004)]. However, these structures are inflexible and suffer from combinatorial explosion in the number of conditional branches as the amount of possible user interaction increases. They also quickly encounter difficulties propagating the consequences of user interaction. Planning, in contrast, offers additional spatial and temporal flexibility. In addition, if an IS system is to allow seamless user interaction, then the overhead of authoring the narrative is potentially far lower with a planning approach than with a branching narrative approach, which requires the enumeration of all conditional branches as it is authored.

- *Story and Discourse*

In narratology, a number of layered models of narrative have been proposed: a two-layer model where narrative has a content layer, “story” and expression layer, “discourse” [Chatman (1978)]; or a three-level model which distinguishes between all of the story events, the “fabula” and those story events presented to the audience, the “sjuzhet” [Rimmon-Kenan (2002)]. Importantly, a plan-based approach provides a consistent technique to manage these different levels of story [Riedl and Young (2010)] and discourse [Jhala and Young (2010)].

- *Boundary Problem*

Finally, we observe that planning provides a means to support user interaction and in so doing provides a solution to the “boundary problem” which arises when all narrative content is authored and user interaction could potentially go beyond this [Magerko (2007b)]. Planning directly sidesteps this problem, as interactions are dealt with at the level of planning, allowing for the generation and re-generation of narrative content based directly on the consequences of user interaction.

Acknowledgements I’d like to thank Fred Charles, David Pizzi, Jonathan Teutenberg, Stephen Gilroy and Marc Cavazza with whom I collaborated on the development of some of the large-scale IS systems discussed in this chapter.

References

E. Amir, A. Chang, Learning partially observable deterministic action models. Journal of

- Artificial Intelligence Research **33**(1), 349–402 (2008)
- M. Arinbjarnar, D. Kudenko, Bayesian Networks: Real-Time Applicable Decision Mechanisms for Intelligent Agents in Interactive Drama, in *Proc. IEEE Conference on Computational Intelligence and Games (CIG)*, 2010
- R. Aylett, J. Dias, A. Paiva, An Affectively Driven Planner for Synthetic Characters, in *Proc. 16th Int. Conference on Automated Planning and Scheduling (ICAPS)*, 2006
- B.-C. Bae, Y.-G. Cheong, R.M. Young, Automated Story Generation with Multiple Internal Focalization, in *Proc. IEEE Conference on Computational Intelligence and Games (CIG)*, 2011, pp. 211–218
- J.A. Baier, F. Bacchus, S.A. McIlraith, A Heuristic Search Approach to Planning with Temporally Extended Preferences, in *Proc. 20th Int. Joint Conference on Artificial Intelligence (IJCAI)*, 2007
- H. Barber, D. Kudenko, Dynamic Generation of Dilemma-Based Interactive Narratives, in *Proc. 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, 2007
- L. Barros, S. Musse, Improving Narrative Consistency in Planning-Based Interactive Storytelling, in *Proc. 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, 2007
- L.M. Barros, S.R. Musse, Towards Consistency in Interactive Storytelling: Tension Arcs and Dead-Ends. *Computers in Entertainment* **6**, 1–17 (2008)
- J. Bates, The role of Emotion in Believable Agents. *Communications of the ACM* **37**, 122–125 (1994)
- J.D. Bolter, R. Grusin, *Remediation: Understanding New Media* (MIT Press, Cambridge (Massachusetts)), 1999
- B. Bonet, H. Geffner, Planning as heuristic search. *Artificial Intelligence* **129**, 5–33 (2001)
- M. Brenner, Creating Dynamic Story Plots with Continual Multiagent Planning, in *Proc. of 24th National Conference on Artificial Intelligence (AAAI)*, 2010
- D. Cage, *Heavy Rain*, Atari, Quantic Dream, 2010
- M. Cavazza, F. Charles, S. Mead, Emergent Situations in Interactive Storytelling, in *Proc. of ACM Symposium on Applied Computing (SAC)*, 2002a
- M. Cavazza, F. Charles, S.J. Mead, Character-based interactive storytelling. *IEEE Intelligent Systems* **17**(4), 17–24 (2002b)
- M. Cavazza, D. Pizzi, F. Charles, T. Vogt, E. André, in *Proc. of 8th Int. Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2009
- A. Champandard, *Behaviour Trees for Next-Gen Game AI*, 2008. <https://aigamedev.com/insider/presentations/behavior-trees/>
- D. Chapman, Planning for conjunctive goals. *Artificial Intelligence* **32**, 333–377 (1987)
- F. Charles, M. Cavazza, Exploring the scalability of character-based storytelling, in *Proc. of 3rd Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2004
- S. Chatman, *Story and Discourse: Narrative Structure in Fiction and Film* (Cornell University Press, Ithaca, New York, 1978)
- Y. Chen, B. Wah, C.W. Hsu, Temporal Planning using Subgoal Partitioning and Resolution in SGPlan. *Journal of Artificial Intelligence Research* **26**, 323–369 (2006)
- A.I. Coles, M. Fox, K. Halsey, D. Long, A. Smith, Managing concurrency in temporal planning using planner-scheduler interaction. *Artificial Intelligence*, 1–44 (2009)
- S. Cresswell, P. Gregory, Generalised Domain Model Acquisition from Action Traces, in *Proc. of 21st Int. Conference on Automated Planning and Scheduling (ICAPS)*, 2011
- F.A.G. Da Silva, A.E.M. Ciarlini, S.W.M. Siqueira, Nondeterministic Planning for Generating Interactive Plots, in *Proc. of 12th Ibero-American Conference on Advances in Artificial Intelligence*, 2010
- D. Dennett, *The Intentional Stance* (MIT Press, Cambridge, MA, 1989)
- M. Fox, D. Long, PDDL2.1: An Extension to PDDL for Expressing Temporal Domains. *Journal of Artificial Intelligence Research* **20**, 61–124 (2003)
- P. Gervás, B. Díaz-Agudo, F. Peinado, R. Hervás, Story plot generation based on CBR. *Knowl.-Based Syst.* **18**(4-5), 235–242 (2005)
- S. Gilroy, J. Porteous, F. Charles, M. Cavazza, E. Soreq, G. Raz, L. Ikar, A. Or-Borichov, U. Ben-Arie, I. Klovatch, T. Hendler, A Brain-Computer Interface to a Plan-based Narrative, in *Proc. of 23rd Int. Joint Conference on Artificial Intelligence (IJCAI)*,

2013

- P. Haslum, Narrative Planning: Compilations to Classical Planning. *Journal of AI Research* **44**, 383–395 (2012)
- R. Hill, J. Gratch, W. Johnson, E. Kyriakakis, C. Labore, R. Lindheim, Toward the Holodeck: Integrating Graphics, Sound, Character and Story, in *Proc. 5th Int. Conference on Autonomous Agents*, 2001
- H. Hoang, S. Lee-Urban, H. Munoz-Avila, Hierarchical Plan Representations for Encoding Strategic Game AI, in *Proc. of 1st Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, 2005
- J. Hoffmann, The Metric-FF Planning System: Translating "Ignoring Delete Lists" to Numeric State Variables. *Journal of Artificial Intelligence Research* **20**, 291–341 (2003)
- J. Hoffmann, Everything You Always Wanted to Know About Planning (but Were Afraid to Ask), in *Proc. 34th Annual German Conference on Artificial Intelligence (KI)*, 2011
- J. Hoffmann, B. Nebel, The FF Planning System: Fast Plan Generation through Heuristic Search. *Journal of AI Research* **14**, 253–302 (2001)
- J. Hoffmann, J. Porteous, L. Sebastia, Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research (JAIR)* **22**, 215–278 (2004)
- ICAPS, *International Planning Competitions*. <http://ipc.icaps-conference.org/>
- IO Interactive, *HITMAN*, Eidos Interactive. Square Enix, 2000
- A. Jhala, R.M. Young, Cinematic Camera Control Using Discourse Planning Techniques, in *Proc. 20th National Conference on Artificial Intelligence (AAAI)*, 2005
- A. Jhala, R.M. Young, Cinematic Visual Discourse: Representation, Generation, and Evaluation. *IEEE Trans. on Computational Intelligence and AI in Games* **2**, 69–81 (2010)
- B. Karlsson, A. Ciarlini, B. Feijó, A. Furtado, Applying a Plan-Recognition/Plan-Generation Paradigm to Interactive Storytelling, in *Proc. of ICAPS Workshop on AI Planning for Computer Games and Synthetic Characters*, 2006
- B. Kartal, J. Koenig, S.J. Guy, User-Driven Narrative Variation in Large Story Domains using Monte Carlo Tree Search, in *Proc. of 13th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2014
- J.P. Kelly, A. Botea, S. Koenig, Planning with Hierarchical Task Networks in Video Games, in *Proc. of ICAPS-07 Workshop on Planning in Games*, 2007
- R.E. Korf, Real-time heuristic search. *Artificial Intelligence* **42**(2-3), 189–211 (1990)
- B. Li, S. Lee-Urban, G. Johnston, M. Riedl, Story Generation with Crowdsourced Plot Graphs, in *Proc. of 27th National Conference on Artificial Intelligence (AAAI)*, 2013
- D. LLansó, M. Gómez-Martín, P. González-Calero, Self-Validated Behaviour Trees through Reflective Components, in *Proc. 5th Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, 2009
- B. Magerko, A Comparative Analysis of Story Representations for Interactive Narrative Systems, in *Proc. 3rd Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, 2007a
- B. Magerko, Evaluating preemptive story direction in the interactive drama architecture. *Journal of Game Development* (2007b)
- M. Mateas, A. Stern, Structuring Content in the Façade Interactive Drama Architecture, in *Proc. of 1st Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2005
- M. Mateas, A. Stern, A behavior language for story-based believable agents. *IEEE Intelligent Systems* **17**(4), 39–47 (2002)
- N. McIntyre, M. Lapata, Learning to Tell Tales: A Data-driven Approach to Story Generation, in *Proc. of 47th Meeting of the Association for Computational Linguistics (ACL)*, 2009
- Merriam-Webster, *Dictionary and Thesaurus Online*, <http://www.dictionaryapi.com/>
- J. Orkin, D.K. Roy, Understanding Speech in Interactive Narratives with Crowdsourced Data, in *Proc. of 8th Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2012
- R. Paul, D.K. Charles, M. McNeill, D.M.G. McSherry, Adaptive Storytelling and Story Repair in a Dynamic Environment, in *Proc. 4th Int. Conference on Interactive Digital Storytelling (ICIDS)*, 2011

-
- A. Piacenza, F. Guerrini, N. Adami, R. Leonardi, J. Teutenberg, J. Porteous, M. Cavazza, Changing Video Arrangement for Constructing Alternative Stories, in *Proc. of 19th ACM Int. Conference on Multimedia*, 2011
- D. Pizzi, M. Cavazza, J.-L. Lugin, Extending character-based storytelling with awareness and feelings, in *Proc. 6th Int. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2007a
- D. Pizzi, F. Charles, J.-L. Lugin, M. Cavazza, Interactive Storytelling with Literary Feelings, in *Proc. 2nd Int. Conference on Affective Computing and Intelligent Interaction (ACII)*, 2007b
- J. Porteous, M. Cavazza, F. Charles, Narrative Generation through Characters' Point of View, in *Proc. of 9th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2010
- J. Porteous, F. Charles, M. Cavazza, NetworkING: using Character Relationships for Interactive Narrative Generation, in *Proc. of 12th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2013
- J. Porteous, J. Teutenberg, F. Charles, M. Cavazza, Controlling Narrative Time in Interactive Storytelling, in *Proc. of 10th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2011a
- J. Porteous, J. Teutenberg, D. Pizzi, M. Cavazza, Visual Programming of Plan Dynamics using Constraints and Landmarks, in *Proc. of 21st Int. Conference on Automated Planning and Scheduling (ICAPS)*, 2011b
- J. Porteous, A. Lindsay, J. Read, M. Truran, M. Cavazza, Automated Extension of Narrative Planning Domains with Antonymic Operators, in *Proc. of 14th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2015
- J. Porteous, M. Cavazza, F. 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)
- M. Riedl, Narrative Generation: Balancing Plot and Character, PhD thesis, Department of Computer Science, North Carolina State University, 2004
- M. Riedl, Incorporating Authorial Intent into Generative Narrative Systems, in *Proc. of AAAI Spring Symposium on Intelligent Narrative Technologies*, 2009
- M. Riedl, A. Stern, Believable Agents and Intelligent Story Adaptation for Interactive Storytelling, in *Proc. of 3rd Int. Conference on Technologies for Interactive Digital Entertainment (TIDSE)*, 2006
- M. Riedl, R.M. Young, Open-world planning for story generation, in *Proc. of 19th Int. Joint Conference on Artificial Intelligence (IJCAI)*, 2005
- M. Riedl, C. Saretto, R.M. Young, Managing interaction between users and agents in a multi-agent storytelling environment, in *Proc. of 2nd Int. Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2003
- M.O. Riedl, R.M. Young, Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research* **39**, 217–267 (2010)
- S. Rimmon-Kenan, *Narrative Fiction: Contemporary Poetics* (Methuen Routledge (2nd Edition), New York, 2002)
- M.L. Ryan, *Avatars of Story* (University of Minnesota Press, USA, 2006)
- N.M. Sgouros, G.K. Papakonstantinou, P. Tsanakas, A Framework for Plot Control in Interactive Story Systems, in *Proc. of 13th National Conference on AI (AAAI)*, 1996
- S. Sina, A. Rosenfeld, S. Kraus, Generating content for scenario-based serious games using CrowdSourcing, in *Proc. of 28th National Conference on Artificial Intelligence (AAAI)*, 2014
- F. Sparacino, Sto(Ry)Chastics: A Bayesian Network Architecture for User Modeling and Computational Storytelling for Interactive Spaces, in *Proc. 5th Int. Conference on Ubiquitous Computing (UbiComp)*, 2003
- U. Spierling, S.A. Weiß, W. Müller, Towards Accessible Authoring Tools for Interactive Storytelling, in *Proc. 3rd Int. Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE)*, 2006
- R. Swanson, A.S. Gordon, A Data-Driven Case-Based Reasoning Approach to Interactive Storytelling, in *Proc. 3rd Joint Int. Conference on Interactive digital storytelling (ICIDS)*, 2010

-
- R. Swanson, A.S. Gordon, Say Anything: Using Textual Case-Based Reasoning to Enable Open-Domain Interactive Storytelling. *ACM Trans. Interact. Intell. Syst.* **2**(3) (2012)
- J. Teutenberg, J. Porteous, Efficient Intent-based Narrative Generation Using Multiple Planning Agents, in *Proc. of 12th Int. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2013
- T. Trabasso, P.W. van den Broek, Causal Thinking and the Representation of Narrative Events. *Journal of Memory and Language* **24**, 612–630 (1985)
- G. Tully, S. Turner, Integrated Decision Points for Interactive Movies, in *Proc. 2nd Int. Conference on Technologies for Interactive Digital Storytelling and Entertainment (TIDSE)*, 2004
- WordNet, *A lexical database for English*, <https://wordnet.princeton.edu/wordnet/>
- R.M. Young, Creating Interactive Narrative Structures: The Potential for AI Approaches, in *AAAI Spring Symposium on Artificial Intelligence and Entertainment*, 2000
- R.M. Young, An Overview of the Mimesis Architecture: Integrating Intelligent Narrative Control into an Existing Gaming Environment, in *AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, 2001
- R.M. Young, Story and Discourse: A Bipartite Model of Narrative Generation in Virtual Worlds. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems* **8** (2007)
- R.M. Young, Notes on the Use of Plan Structures in the Creation of Interactive Plot, in *AAAI Fall Symposium on Narrative Intelligence*, 1999
- H.H. Zhuo, S. Kambhampati, Action-model Acquisition from Noisy Plan Traces, in *Proc. of 23rd Int. Joint Conference on Artificial Intelligence (IJCAI)*, 2013
- H.H. Zhuo, Q. Yang, D.H. Hu, L. Li, Learning complex action models with quantifiers and logical implications. *Artificial Intelligence* **174**(18), 1540–1569 (2010)