

Planning Formalisms and Authoring in Interactive Storytelling

Fred Charles¹, Miguel Lozano², Steven J. Mead¹, Alicia Fornes Bisquerra²
and Marc Cavazza¹

¹ School of Computing and Mathematics, University of Teesside,
UK-TS1 3BA, Middlesbrough, United Kingdom
{f.charles, steven.j.mead, m.o.cavazza}@tees.ac.uk
²University of Valencia, Dr Moliner 50, (Burjassot), Valencia, Spain
{miguel.Lozano, Alicia.f.bisquerra}@uv.es

Abstract. Interacting Storytelling systems integrate AI techniques such as planning with narrative representations to generate stories. In this paper, we discuss the use of planning formalisms in Interactive Storytelling from the perspective of story generation and authoring. We compare two different planning formalisms, Hierarchical Task Network (HTN) planning and Heuristic Search Planning (HSP). While HTN provide a strong basis for narrative coherence in the context of interactivity, HSP offer additional flexibility and the generation of stories and the mechanisms for generating comic situations.

1 Introduction

Research in Interactive Storytelling (IS) has developed in a spectacular fashion over the last few years, due to progress in the integration of real-time Artificial Intelligence techniques into graphic environments, as well as a growing interest in its potential applications in training or entertainment.

The diversity of approaches is not easy to capture into a simple classification, yet a small set of key problems has been identified by early IS research: the trade-off between interactivity and storytelling, the duality between character and plot, narrative causality, the problem of narrative control, and the relations between story generation and presentation [13] [14] [24] [25].

With the exception of emergent storytelling, IS systems rely on various AI techniques to support their behaviour, from Assumption-based Truth Maintenance System (ATMS) [20] and logic programming [8] to planning systems [23] [24]. The knowledge content consists in narrative formalisations often inspired from narrative theories, including those of Aristotle [15], Propp [8] [17], Barthes [2] [5], Bremond [22], etc. These narrative theories try to provide input on the dynamics of the story and its dramatic nature [14], on the formalisation of narrative actions, or on that of actors' roles. The common denominator of story formalisation, however simplistic it might first appear, has been to consider a story as a sequence of actions related through some form of causality. Some research has focussed on causal representations, e.g. the use

of ATMS by Sgouros et al [20] or StoryNets by Swartout et al. [21], while others have concentrated on the generation and control of a sequence of narrative actions.

However, if we restrict ourselves to the structural aspects of a narrative, seen in the abstract as a sequence of actions, it does not come as a surprise that AI planning formalisms have been one of the main techniques supporting interactive storytelling [13] [21] [24], as planning precisely deals with the generation of action sequences fulfilling a goal. Besides, planning is also a generic description for an agent's behaviour [7] and in that respect, can also be used to represent the role of a character in a character-based approach [4].

In addition, the use of planning techniques in storytelling is independent of the paradigm selected, whether plot-based or character-based. This can be interpreted as a consequence of the duality exhibited by some narrative representations. If traditional narrative functions are extended to re-introduce agency, then an equivalence can be established between a plot representation based on these narrative functions and a collection of well-identified roles for the story characters. This can be illustrated by considering some traditional narrative functions and relating them to roles: for instance, the role of influencer, as introduced by Bremond, is naturally echoing many of Propp's functions.

2 Planning in Character-based Storytelling

Several different planning techniques have been reported to be used in IS research and related behavioural animation research, for instance POP-style planning [12] [23]. The rationale for the use of these techniques has been in many cases just their availability as part of the background work of interdisciplinary teams assembled to work on IS projects. On the other hand, potential uses of planning have been clearly identified, such as narrative control, supporting suspense [23] and assisting authoring [19].

While planning has been mostly used to formalise the plot itself, in the remainder of the paper we will describe our own approach in which planning is used to describe individual characters' roles, the plot being dynamically produced through character interaction. We will illustrate the discussion with our virtual sitcom system, which is inspired from the set of characters and situations in the popular Friends™ series. The system has been fully implemented using the 3D game engine Unreal Tournament™ as a graphic environment, into which two different planning modules have been integrated.

Using planning techniques within a character-based approach [4], we are essentially concerned with the following aspects of planning: i) the relation between planning and the specific authoring process for interactive narratives ii) how planning would model relevant narrative phenomena, such as beats, suspense or even the comic nature of situations and iii) how planning supports story variation, i.e. the change in the course of action.

2.1 Planning and Narrative Representations

When using plans to represent characters' roles, we are working under the following assumption: the actor's behaviour in the story will correspond to an actual plan, whose various stages correspond to narrative actions. This means that, from a diegetic perspective, the actions witnessed by the spectator actually take the shape of a plan (e.g. to rob a bank, seduce a lover, etc.). From a representational perspective, this consists in: i) distributing the various narrative functions which represent the storyline into each actor taking a role in these functions and ii) instantiating the plan of each actor in terms of actual narrative events. The latter point ensures that the behaviour of the actor is actually perceived as a plan as well.

We now have to discuss how planning can generate events of narrative and dramatic interest. If the plan is to be perceived as such by the spectator, this means that elements of the planning process itself can be given an interpretation in terms of dramatic value. The first element to be considered is goal-orientation. As with traditional narratology, this is a condition for the believability of character behaviour. In line with what was previously said, even in a character-based approach, key elements of the storyline will appear as the main goals of a feature characters. Traditional narrative elements such as antagonism [18] have a translation in planning terms as well: for instance, two rival characters will have competing plans for the same object. When one character has a model of another character's plan and specifically tries to contrast it, the term of counter-planning is used instead.

Traditional models in narratology, whether Aristotle's or Propp's, describe a course of action quite incompatible with the smooth unfolding of the main character's plan. This means that plan failure is an essential aspect of narrative interest. This happens to be the case across story genres and, while the overall goal can be eventually achieved, dramatic interest can be obtained by "local" plan failure and subsequent re-planning or plan repair. A key aspect in the use of planning formalisms in storytelling thus consists in their ability to support re-planning and, to some extent, to offer representations embedding the potential for failure. In the next sections, we will be discussing how this can be supported by two different planning formalisms based on our own experience in developing IS prototypes.

3. HTN Planning for Interactive Storytelling

In the first instance, we required a formalism with good representational abilities that could still support interleaving planning and execution, which is a strong requirement as the characters have to generate their actions in real-time within an environment where they interact with each other. HTN planning is recognised as being appropriate to knowledge-intensive domains such as IS, and this was the initial rationale for selecting it. Overall, HTN planning consists in using a hierarchical task model (represented as an AND/OR graph), in which each task is decomposed into sub-tasks until they can be described in terms of primitive actions (in IS, these actions are those actually played in the story). The HTN can then be searched to extract a task decomposi-

tion corresponding to a solution plan. In our case, we adapted the search mechanism so as to obtain a real-time version interleaving planning and execution. HTN implements forward search from the goal state which brings the unique property that during planning itself the state of the world is known at all times. We have given several descriptions of our HTN planner [2] [3] [4] and will rather concentrate in the next sections on the representational aspects.




Fig. 1. HTN representation of lead character behaviour.

Authoring an HTN can be seen as the description of a character's role in a traditional sense. Its actions have to be decomposed into lower-level elements, until they can be described as elementary physical interventions on stage. From a narrative perspective, the natural hierarchy of an HTN can thus be used to describe several levels of representations, helping to attribute a narrative interpretation to the HTN nodes. The top-level node will correspond to the story goal, the first layer nodes the main temporal decomposition (scenes) and lower-level nodes to various tactics used in pursuing the scene goals. This is how in our Friends™ scenario the top-level goal (to invite Rachel out (see Figure 1)) corresponds to the baseline story, while the first level of decomposition describes the various steps required to achieve that objective (such as learning about her taste, finding a way to talk to her in private, etc.). Further decompositions correspond to the various courses of action to achieve a given sub-goal: for instance, learning about Rachel's preferences can be achieved by observing her, asking her friends, etc. (see Figure 2).




Fig. 2. Illustration of story instantiation.

As representations, HTN rest on two main hypotheses, which are task decomposability and total ordering. Task decomposability assumes that the action can be split into sub-tasks with no interaction between them, neither positive nor negative. Negative interactions happen when the outcome of a sub-task undoes what has been previously achieved by another one: for instance, Ross might prepare himself for going out but if he fights with Joey he will have to get dressed again. Positive interactions are described when carrying out a given task facilitates the execution of a subsequent one. For instance, if Ross decides to clean the flat and in order to do so asks everybody to leave, this will facilitate talking to Rachel in private later on.

The HTN representation offers a good compromise between authoring and the provision of story variability within a baseline story. The reason is that total ordering supports the decomposition into scenes, preserving a basic temporal structure for the high-level actions, while offering variability in terms of action selection at lower-levels. This variability is based on various models of action preferences based on the emotional status of characters and the need to re-plan when a previously selected action fails.

The total ordering assumption may preserve the decomposition of a story into scenes, but such decomposition at every level can be too restrictive. In particular, the absence of backward effects, undoing previous tasks, is a limitation for narrative representation. In our use of HTN, actions can fail, in which case the HTN planner provided facilities for backtracking within the sub-task governing that action. However, once a task had been completed, there was no way of undoing the outcome of this task, save re-starting the whole story, which is rarely acceptable. Also, re-planning tends to take place at the level immediately superior in the task decomposition, which generally rules out dramatic changes of strategy. In addition, HTN can become complex and difficult to manage, though on the other hand they provide a principle for surveyability and global authoring, which is clearly relevant in IS.

4. Heuristic Search Planning and Story Generation

We considered the limitations of the HTN approach, particularly for the generation of more complex and intricate stories enabling the emergence of long-distance dependencies. This suggested using a more modular description for the individual narrative actions that could incorporate various forms of long-distance interactions. A traditional STRIPS-like representation [6] as used in planning, appeared appropriate to this objective, as it would enable the description of individual actions together with the changes they bring, both positive and negative, to the story (see Figure 3).




Fig. 3. HSP system overview.

A formal equivalence has been established between STRIPS formalisms and HTN in terms of their expressivity. However, as it is often the case, that formal equivalence does not really impact on the knowledge acquisition phase, i.e. on the procedures by which formalisms are instantiated with narrative content by users/designers and in practice, the features in our two systems are quite different.

STRIPS formalisations are compatible with a variety of planning approaches: we were in search of one that could be computationally efficient and make direct use of the STRIPS formulations so as not to interfere with authoring in the first instance.

Heuristic Search Planning (HSP) is a state-of-the-art planning technique [1] developed recently, which has immediately demonstrated great potential, outperforming many other planning techniques when compared to them on traditional planning benchmark problems.

HSP uses a STRIPS-based representation for problem description and searches the space of states from the initial state, using a traditional heuristic search algorithm and a heuristic automatically extracted from the STRIPS formulation. In simple terms, this heuristic measures the prospect of reaching the goal state from each of the new generated states, by computing the length of an operator chain that would generate the propositions composing the goal state. The heuristic function works on a relaxed prob-

lem, ignoring delete-lists of operators and only considering facts in their add-list. This heuristic approximation ignores positive interactions among sub-goals that can make one goal simpler after a second one has been achieved. And of course, negative interaction between sub-goals is ignored as delete lists are not taken into account.

In our implementation, the search algorithm is based on the MinMin algorithm of Pemberton and Korf [16], as it is required to interleave planning and execution [11]. The heuristic is computed using the method described in [10].

We have thus experimented with the same “Friends™” scenario using HSP to power individual characters. For each character, we defined a set of operators corresponding to the behaviour of that agent; this time operators can have a non-empty delete lists and support long-term dependencies. Figure 4 shows two typical STRIPS operators. This second planner works by selecting the most appropriate operator at each step, using the formal HSP heuristic following the end goal as a narrative drive. At each step, the totality of applicable operators are considered, which supports more dramatic re-planning in an actor’s plan than offered with HTN re-planning, which often involves re-planning from the immediately superior level. For instance, while in an HTN failure to learn about Rachel’s preferences by asking her will lead e.g. to asking her friend to solve the same sub-task, in HSP a new course of action can be started, such as going out to the shops, from where another solution can be found. While HSP supports more variability, narrative consistency is maintained by the existence of a single narrative drive and the nature of the heuristic.

```
(def-operator: use-rachels-pda
  (make-operator
    :pre-conditions `(:need-gift-idea)
    :exe-condition :pda-free
    :add-list `(:info-gift-flowers)
    :delete-list `(:need-gift-idea))

(def-operator: purchase-flowers
  (make-operator
    :pre-conditions `(:info-gift-flowers :at-shop :has-money)
    :exe-condition :flowers-available
    :add-list `(:has-flowers)
    :delete-list `(:has-money)))
```

Fig. 4. Example of STRIPS operators defined for the Friends™ scenario.

Not only does HSP offer more variability in the action sequences, but by enabling long-distance dependencies, it supports the generation of comic situations in which failure does not only result from external intervention but from the poor design of the character’s plan. For instance, in Figure 5 Ross decides to buy Rachel flowers at an early stage of the story. He cannot phone to book the restaurant from the flat, nor can he leave the roses unattended. He thus carries the roses with him, but meets Phoebe who thinks that the roses are for her.




Fig. 5. Example of story variability.

5. Discussion

We can now revisit the fundamental relations between planning and narrative interest introduced in section 4 from the HSP perspective. With HTN planning, the character is executing a pre-defined and fundamentally correct plan. Pre-defined here does not mean that the plan is compiled but that its higher-level goals are fixed, flexibility and generativity taking place at the lower levels. As discussed in section 2, narrative interest comes from action failure within this pre-defined plan, either due to user intervention or interference by other agents [3].

Conversely, with HSP planning, the agent is generating its own plan from a single narrative drive equivalent to the HTN's top-level goal. Under these circumstances narrative interest can derive, not only from external interaction but also from the qualities of the plan assembled. Basically, misconceptions or blunders that derive from the nature of the plan itself will produce comic situations. Very often this will be due to the ignorance of long-term dependencies in the actions carried out. This emphasises the need for a formalism actually supporting this kind of dependencies.

In terms of generativity, the main mechanism in HTN is the interaction between characters' plans at the level of terminal actions. However, this form of generation produces instances of plots within a given storyline as defined by the role HTN for the main characters. This, in agreement with our initial objectives that were to allow user intervention while still maintaining the basic storyline, cannot be used to explore a generation of larger set of stories.

On the other hand, HSP planning will produce as many different stories as potentially conceivable from the initial narrative drive, using a whole population of operators. It favours a more exploratory approach, which is well suited to experiment with the mechanisms for generating comic situations.

In HTN, links between actions are simpler to describe, as the detailed consequences of actions in terms of state do not have to be described. On the other hand, the detailed description of the various constraints in STRIPS operators is also the opportunity to generate complex situations as shown above.

Conclusions

We have discussed the role of planning techniques in the implementation of IS systems and the associated problems of authoring and representation. By comparing two different planning formalisms, HTN planning and HSP, we have illustrated from yet another perspective the trade-off between narrative control and story generation. HTN offer clear authoring principles and a global vision on the baseline plot, while HSP offer greater flexibility in the definition of action and more variability in the stories generated. In particular, support of the generation of dramatic situations as a property of the plan itself rather than as a consequence of external intervention.

The individual definition of STRIPS operators in authoring the character's actions is highly flexible but does not provide any basis for narrative integration. In principle, this could constitute a limitation when interacting with professional scriptwriters. We are still at an early stage of our experiments with this technique, but it might be the case that its potential use lies more in the exploration of potential storylines, taking advantage of its generative capabilities.

References

1. Bonet, B. and Geffner, H. Planning as Heuristic Search. *Artificial Intelligence Special Issue on Heuristic Search*, 129(1), pp. 5-33, 2001.
2. Cavazza, M., Charles, F. and Mead, S.J. Characters in Search of an Author: AI-based Virtual Storytelling. *Proceedings of the First International Conference on Virtual Storytelling (ICVS 2001)*, Avignon, France, Lecture Notes in Computer Science, n. 2197, Springer Verlag, pp. 145-154, 2001.
3. Cavazza, M., Charles, F. and Mead, S.J. Interacting with Virtual Characters in Interactive Storytelling. *ACM Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, pp. 318-325, 2002.
4. Cavazza, M., Charles, F. and Mead, S.J. Character-based Interactive Storytelling, *IEEE Intelligent Systems*, special issue on AI in Interactive Entertainment, pp. 17-24, 2002.
5. Douglas J., Gratch, J. Interactive Storytelling, or How AI, Hollywood, and Multiprocessing Operating Systems can Live Happily Ever After", *Proceedings of the First International Conference on Virtual Storytelling (ICVS 2001)*, Avignon, France, Lecture Notes in Computer Science, n. 2197, Springer Verlag, 2001.
6. Fykes R. and Nilsson N., STRIPS: A new approach to the application of theorem proving to problem solving *Artificial Intelligence*, 2:189-208, 1971.
7. Geib, C. and Webber, B.: A consequence of incorporating intentions in means-end planning. Working Notes – *AAAI Spring Symposium Series: Foundations of Automatic Planning: The Classical Approach and Beyond*. AAAI Press, 1993.
8. Grasbon, D. and Braun, N. A Morphological Approach to Interactive Storytelling. *Proceedings of on Artificial Intelligence and Interactive Entertainment, Cast'01, Living in Mixed Realities*, Sankt Augustin, Germany, 2001.
9. Gratch J. and Marsella S. Tears and Fears: Modeling emotions and emotional behaviors in synthetic agents, in *Proceedings of the 5th International Conference on Autonomous Agents*, Montreal, Canada, June 2001.

10. Liu Y., Koenig S. and Furcy D., Speeding Up the Calculation of Heuristics for Heuristic Search-Based Planning, *In Proceedings of the National Conference on Artificial Intelligence*, pp. 484-491, 2002.
11. Lozano, M., Mead, S.J., Cavazza, M. and Charles, F. Search Based Planning: A Model for Character Behaviour, in *Proceedings of the 3rd on Intelligent Games and Simulation, GameOn-2002*, London, UK, November 2002.
12. Magerko, B. A Proposal for an Interactive Drama Architecture, *AAAI 2002 Spring Symposium Series: Artificial Intelligence and Interactive Entertainment*, March 2002.
13. Mateas, M.: An Oz-Centric Review of Interactive Drama and Believable Agents. *Technical Report CMU-CS-97-156*, Department of Computer Science, Carnegie Mellon University, Pittsburgh, USA, 1997.
14. Mateas, M. Narrative Intelligence: An Introduction to the NI Symposium. In M. Mateas and P. Sengers (Eds.), *Working notes of the Narrative Intelligence Symposium*, AAAI Fall Symposium Series. Menlo Park: Calif.: AAAI Press, 1999.
15. Mateas, M.: A Neo-Aristotelian Theory of Interactive Drama. *Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, AAAI Press 2000.
16. Pemberton, J.C. and Korf, R.E.: Incremental Search Algorithms for Real-Time Decision Making. *Proceedings of the Second Artificial Intelligence Planning Systems Conference (AIPS-94)*, pp. 140-145, 1994.
17. Prada, R., Machado, I., and Paiva, A.: Teatrix: Virtual Environments for Story Creation. *Proceedings of Intelligent Tutoring Systems 2000*, Canada, 2000.
18. Propp, V.,. *Morphology of the Folktale*. University of Texas Press: Austin and London, 1968.
19. P. Rizzo, M.V. Veloso, M. Miceli, and A. Cesta, Goal-based personalities and social behaviors in believable agents, *Applied Artificial Intelligence*, 13, pp. 239-27, 1999.
20. Sgouros, N.M., Papakonstantinou, G. and Tsanakas, P., 1996. A Framework for Plot Control in Interactive Story Systems, *Proceedings AAAI'96*, Portland, AAAI Press, 1996.
21. Swartout, W., Hill, R., Gratch, J., Johnson, W.L., Kyriakakis, C., LaBore, C., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thiebaut, M., Tuch, L., Whitney, R. and Douglas, J. Toward the Holodeck: Integrating Graphics, Sound, Character and Story. in *Proceedings of the Autonomous Agents 2001 Conference*, 2001.
22. Szilas, N.: Interactive Drama on Computer: Beyond Linear Narrative. Papers from the AAAI Fall Symposium on Narrative Intelligence, *Technical Report FS-99-01*, AAAI Press, 1999.
23. Young, R.M. Cognitive and computational models of suspense: Towards the automatic creation of suspense in interactive narratives, talk presented at Interactive Frictions: The Conference Produced at the Pressure Point Between Theory and Practice: An International Conference on Interactive Narrative Friday, June 4, 1999, University of Southern California, Los Angeles, CA, 1999.
24. Young, R.M. Creating Interactive Narrative Structures: The Potential for AI Approaches. *AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment*, AAAI Press, 2000.
25. Young, R.M.: An Overview of the Mimesis Architecture: Integrating Intelligent Narrative Control into an Existing Gaming Environment. *Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, AAAI Press 2001.