Characters in Search of an Author: AI-Based Virtual Storytelling

Marc Cavazza[1], Fred Charles[1], and Steven J. Mead[1]

[1] School of Computing and Mathematics, University of Teesside,
UK-TS1 3BA, Middlesbrough, United Kingdom
{m.o.cavazza, f.charles, steven.j.mead}@tees.ac.uk

Abstract. In this paper, we present the first results obtained with an interactive storytelling prototype. Our main objective is to develop flexible character-based systems, which nevertheless rely on narrative formalisms and representations. Characters' behaviours are generated from plan-based representations, whose content is derived from narrative formalisms. We suggest that search based planning can satisfy the real-time requirements of interactive storytelling, while still being compatible with the narrative formalisation we are pursuing. We then describe into greater detail a short episode generated by the system, which illustrates both high-level results and technical aspects, such as re-planning and user intervention. Further work will be dedicated to developing more complex narrative representations and investigating the relations between natural language semantics and narrative structures in the context of interactive storytelling.

E dov'è il copione?
— È in noi, signore.

Luigi Pirandello, *Sei personnagi in cerca d'autore*

1 Introduction

In this paper, we describe the principles behind a virtual interactive storytelling prototype, in which a generic storyline played by artificial actors can be influenced by user intervention. The final applications we are addressing consist in being able to alter the ending of stories that have an otherwise well-defined narrative structure, thus reconciling interactivity with story authoring. Ideally, this would make possible to alter the otherwise dramatic ending of a classical play towards a merrier conclusion.

There has been much recent work in interactive storytelling that has developed a wide range of perspectives; emergent storytelling [1] [2], user-centred plot resolution [3], character-based approaches [4] [5], anytime interaction [6] and the role of narrative formalisms [7]. This work has identified relevant dimensions and key problems for the implementation of interactive storytelling, among which: the status of the user, the level of explicit narrative representation and narrative control, the modes of user intervention, and, most importantly, the relations between characters and plot.

Some of the above problems derive from the inherent tension between interaction and narrative [4] [8]. Interactive systems demand user involvement but often at the expense of a real storyline; on the other hand, a strong narrative dimension is traditionally conceived of with a user as spectator rather than actively involved. Our own solution to this problem consists (in accordance with our final objectives stated above) in limiting the user involvement in the story, though interaction should be allowed at anytime. This is achieved by driving the plot with autonomous characters' behaviours, and allowing the user to interfere with the characters' plans. The user can interact either by physical intervention on the set or by passing information to the actors (e.g., through speech input). In this context, the most important aspect of interactive storytelling is the relation between characters and plot. In his now classic play, Pirandello imagined that characters could be collectively in possession of the plot [9]. This is the best illustration, in modern times, of the duality between character and plot, much debated since its introduction by Aristotle [10]. In the next sections, we develop the hypothesis that narrative functions describing a story can be used to generate plan-based behaviours for the characters. Further, we propose that the respective roles for the various characters of a story should be defined from high-level narrative principles, in a similar fashion.

A first interactive storytelling prototype has been fully implemented and runs in a real-time interactive 3D environment [11]. Graphic rendering, character animation and user interaction in the system are based on the Unreal Tournament™ game engine. This engine provides high-quality display at a constant frame rate, while also serving as a software development environment [12]. Besides embedding its own scripting language (UnrealScript™), it can integrate complete C++ modules or communicate via sockets with external software modules. This prototype has been developed in C++ and UnrealScript™, and runs a simple scenario that we describe in the next sections, together with some of the results obtained.

2 Characters-driven Storytelling and Narrative Formalisms

The storyline for our experiments is based on a simple sitcom-like scenario, where the main character ("Ross") wants to invite the female character ("Rachel") out on a date. This scenario comprises a principal narrative element (will he succeed?) as well as situational elements (the actual episodes of this overall plan that can have dramatic significance, e.g., how he will manage to talk to her in private if she is busy, etc.). For this reason, sitcoms appeared as an interesting narrative genre to investigate. Our system is driven by characters' behaviours, represented as plans.

The behaviour of our artificial characters is based on AI planning techniques, as introduced by Webber et al. [13] for high-level behaviours of virtual actors and by Young in the specific case of storytelling [4]. Our perspective on the contents of the character's plans is clearly narrative rather than cognitive.

We are endeavouring to define a proper representational content for narratives within the implementation framework of AI planning representations. This objective can be illustrated by analogy with computational linguistics: linguistic formalisms can be used to analyse natural language through parsing, and the same syntactic descriptions can serve to generate sentences and text. In a similar fashion, we suggest that a true "computational narratology" could be created using the formalisms developed by narrative analysis. These formalisms would serve as a basis for narrative representations from which stories would be generated. Let us consider first that the initial step in formalizing a plan is to describe a Hierarchical Task Network (HTN), i.e. a hierarchy of sub-goals and actions corresponding to the overall plan. HTNs will thus be our target representations for characters' plans: section 3 will describe the actual generation of behaviours from HTNs.

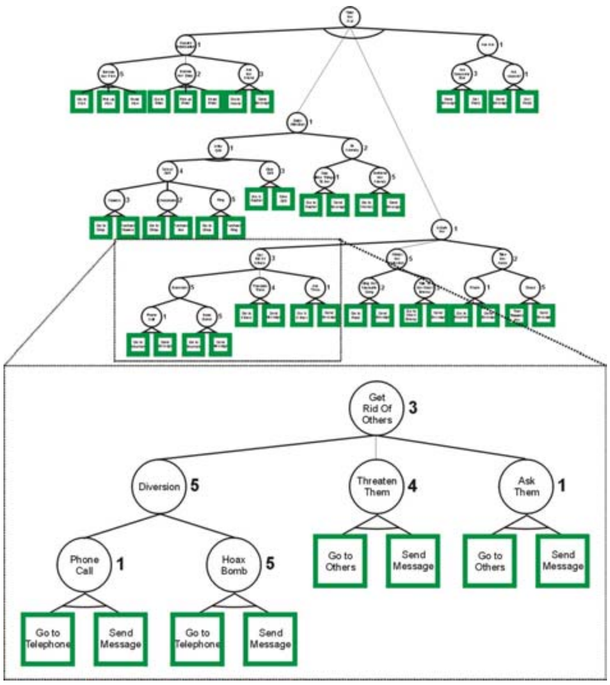

Fig. 1. Ross' Plan

Most work in interactive storytelling has made some reference to narratology. Mateas [8] proposed a neo-Aristotelian framework for interactive stories, Prada et al. [14] referred to Propp's narrative functions as a relevance model for her story-like situations, Szilas [7] has advocated explicit narrative representations based on more recent narrative theories, like those of structuralist authors such as Greimas and Bremond, who have extended narrative formalisms beyond Propp's functions.

A natural approach is to investigate the kind of formalisation attempted in narratology and to determine whether it can be made more computational. Indeed, two sorts of tree-based representations have been introduced in structural narratology, mainly by Barthes: the stemma-like representation [15] and the *proairetic* tree [16]. The former illustrates the decomposition of narrative functions into temporal sequences of lower-level functions. This can be illustrated by considering the overall plan for the Ross character (Figure 1). In order to invite Rachel, he must for instance acquire information on her preferences, befriend her, find a way to talk to her in private, and finally formulate his request (or having someone acting on his behalf, etc.). The latter makes explicit the choices that a character can make at various points[1], i.e. the alternative actions he can take. Dynamic choice of an actual course of action is actually the basis for plot instantiation in interactive storytelling, as otherwise suggested by Young [4]. A given plot will correspond to one and one choice only, together with its long-term consequences. In our scenario, Ross can choose to isolate Rachel from her friends by attracting her attention or by rudely interrupt the conversation. These options (among others) have obviously quite different consequences.

To summarise, we can say that a narrative representation would be an HTN, whose nodes are constituted by various levels of narrative functions, the relationships between the various levels representing composition or alternatives. The HTN represents more than the “role” of the character, as it encompasses potential all variants, at every level. The final level of the HTN consists in terminal actions, i.e. those actions actually played by the character on the virtual stage. A narrative function can correspond to sub-goals at different levels of hierarchy (acquire-information, isolate-her, offer-a-gift): the important point is that the predicative structure of narrative functions (i.e., the other actors involved) is deferred to the lowest compatible level of the hierarchy. This predicative structure is also a function of the main character from which perspective the HTN is described. The initial storyline should actually determine not only the main character plan, but those of other characters as well. This separate definition of roles shall serve as a basis for the dynamic generation of story variants, as individual characters’ plans will interfere with one another, depending on initial conditions and pseudo-random factors. The problem of dependencies between characters’ roles has been described within modern narratology, though not to a formal level. Narrative functions can be refined into bipolar relations between couple of actors, emphasising the asymmetry in their roles [15]. We have adopted this framework to define the respective behaviours of our two leading characters. We started with the overall narrative properties imposed by the story genre; sitcoms offer a light perspective on the difficulties of romance: the female character is often not aware of the feelings of the male character. In terms of behaviour definition, this amount to defining an “active” plan for the Ross character (oriented towards inviting Rachel) and a generic pattern of behaviour for Rachel (her day-to-day activities, subject to variations according to her mood and feelings). This is illustrated on Figure 2. There is significant evidence from narratology studies in favour of this approach: what we have described here is very similar to the respective roles of the main characters in Balzac’s novel *Sarrasine*, which has been entirely analysed by Barthes in his seminal book, “S/Z” [16].

We can now propose a very preliminary methodology for the definition of roles within an overall storyline:

- identify the various roles and the main feature characters
- describe the roles for these main characters as generic plans. In doing so, the predicative structure of narrative functions is refined: narrative functions in

[1] Barthes uses the concept of *proairesis*, or choice between various courses of action, with reference to Aristotle [16].

corporate reference to other characters in their definition: ask-*her*-out(Rachel), ask-*her*-friend(Pheobe), etc.

- enhance the role of feature characters with *proairetic* variations at the appropriate level of description

Fig. 2. Ross' “active” plan and Rachel's generic pattern of behaviours

Another important concept in interactive storytelling is causality [10], as it supports the consequences of interaction, whether it be agent-agent interaction or user intervention. Some interactive storytelling systems make causality explicit in their representations, for instance by using an ATMS [3]. However, in a task network representation based on actions and sub-goals, causality is not explicitly represented. One form of implicit causality is the enabling of further actions by their predecessors in the HTN ordering, but it is not related to interaction and dynamic generation. Other forms of causality are implicit as well: for instance, if, when attempting to talk to Rachel in private, Ross behaves rudely with Pheobe, he might actually upset Rachel and cause her to change feelings in his regard (see the example of section 4). This point illustrates an important practical equivalence between choice and causality, which has been described by Raskin [10]. In a plot-based approach [3] causality can be explicitly represented: in a character based approach, the proairetic aspect is dominant. The character-plot duality has thus a translation in terms of causal representations.

3 AI Planning for Characters’ Behaviour

AI planning is used to implement characters’ behaviour. The planning mechanism should produce a real-time plan from the plan-based narrative representations. An essential requirement, common to all virtual actors evolving in dynamic environments, is that planning and execution should be interleaved [17] [18]. A specific constraint of interactive storytelling is that actions executed by the characters should be properly played in the context of the story: we call this aspect, which concerns the visual presentation of action, *dramatisation*. This is an important aspect, as the user will determine his intervention (if any), according to the meaning he attributes to the characters’ actions. Finally, the approach should support re-planning when the initial plan fails, due to interference from other characters or the user.

The planning system generates a plan for each character, using the HTN describing its own role. From a formal perspective, if we assume that the various sub-goals are independent, planning can be directly achieved by searching the AND/OR graph corresponding to the HTN [19]. In the generic case, the HTN would have to undergo a complex linearisation process beforehand, but in the case of goal independence the solution plan is a direct sub-graph of the HTN[2]. The system can thus use an algorithm such as AO* to produce a solution sub-graph, whose terminal nodes form a sequence of actions to be played by the virtual actor [11]. The standard AO* algorithm comprises a forward expansion, which generates a solution basis (the best solution sub-graph), and a backwards propagation from terminal nodes, which updates the value of the solution expanded. However, AO* alone does not meet our requirements for planning. We have thus developed a "real-time" variant of AO* that interleaves planning and execution. Our planner uses left-to-right depth-first search in a similar fashion than the MinMin algorithm of Pemberton and Korf [20]. It plans forward, until it reaches the first activable terminal action[3], which is then carried out by the character and appropriately dramatised in the virtual environment. The outcome of action execution can be propagated back into the search process by taking advantage of the rollback mechanism which is part of AO*, which is essentially deferred to action execution in our real-time variant. This supports re-planning on action failure, which is one of the mechanisms for story generation. For instance, Ross wanted to know Rachel better by reading her diary. But the user has hidden the diary in a safe place (Figure 3). Ross started to execute his plan, but only realised the diary was missing when reaching its default location. Ross has to find a new way of acquiring information on Rachel, and re-plans a solution for that sub-goal, which consists in asking her friend Phoebe. He starts the execution of this new sequence by looking for Phoebe.

Top-down planning alone is not sufficient to cope with the executability conditions of actions in a dynamic environment. For instance, Ross might not want to steal Rachel's diary if he can be spotted by Monica. To solve this kind of problem, Geib and Webber have proposed to complement top-down planning with situated reasoning [21] [22]. This can be successfully applied in conjunction with real-time planning: the main plan can be interrupted to cope with situated actions and control later returned to the plan after updating the action pre-conditions. For instance, if at an early stage of his plan, Ross bumps into Rachel in a corridor, he cannot just ignore her, but this situation cannot be incorporated into the top-down plan: it has to be treated specifically. Situated reasoning can take place at various levels of the plan hierarchy. For instance, it can enforce generic "priority" behaviours, like avoiding Rachel at the early stages of the plan. At the lowest level of the plan, situated reasoning can even constitute an alternative to re-planning. For instance, if Ross wants to read Rachel's diary but she is using it, he can wait for her to finish (unnoticed by her), rather than find another

[2] To a large extent, sub-goal independence appears to be a property of narrative representations, though this point deserves further investigation.

[3] A similar approach, has been described by Geib [21] as part of his incremental planner "It-Plans". However, our system proceeds depth-first towards the first executable action.

source of information. This form of situated reasoning is based on the duration of actions and the nature of their resources.

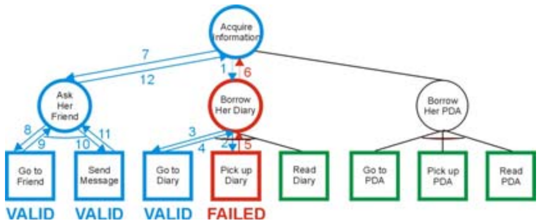

Fig. 3. Re-planning following user intervention

The final story instantiation is mostly determined by the interaction between characters, i.e. by how their plans eventually result in joint action. At the algorithmic level (i.e., RTAO*), there are no intrinsic synchronisation mechanisms between the search processes that drive the two character's behaviours. Rather, using similar principles to those governing user intervention, the characters can interact via their physical environment, by competing for resources for action (narrative objects or other characters). For instance, Ross can influence Rachel's activity and make her more available by e.g., taking care of one of her duties without telling her. Or he can look for some information from Phoebe, but she left to do some shopping with Monica. It is this interaction between the two behaviours that produces much of the situational elements of the story (apart from the final conclusion). To some extent the actual plot can be seen as the "cross product" of the characters plans[4]. There is a number of factors that concurs to make the plot not predictable from the user's perspective, mostly related to actions' duration and competition for action resources. For instance, depending on their initial random positions, some actors can engage in a conversation and become unavailable to others. Similarly, they can be first to reach some objects of narrative importance (telephone, diary, etc.), which will cause other characters' to change their plans, creating new interactions, etc. However, the important point is that characters always keep track of their long-term goals, which differentiates interactive storytelling from simulation-based computer games.

Finally, user intervention is another source of plan variability. The user can interfere with either the execution of the plans' terminal actions or with the plan goals: this determines two major modes of interaction: physical interaction and linguistic interaction [23]. Physical interaction takes place when the user interferes with plans resources on the virtual set, for instance by stealing an object that the artificial actor might use to reach its goal (the "diary" example above). Linguistic interaction is based on speech input that directly passes information to the artificial actors, altering their intentions and goals. For instance, the user can issue a recommendation such as "try to be nice", using speech recognition. This should rule out any rude behaviour towards Rachel or her friends, such as sending her friends away to talk to her. Once again, the effects of such a "doctrine statement" [13] can be implemented by revising the heuristic values attached to certain node categories in the plan graph.

---

[4] This interesting metaphor was suggested to us by an anonymous reviewer.

4 Dynamic Story Generation: First Results

In this section, we describe into further detail a complete example obtained from the prototype. This example will illustrate user intervention, re-planning and the use of moods to propagate causality between various characters.

In order to get the information he needs, Ross goes to read Rachel's diary (a). However, he realises that somebody moved the diary (b). So instead, he decides to meet Phoebe to ask her about Rachel (c). In the meantime, Rachel is talking to Monica (d). In order to talk privately with Rachel, Ross is ordering Monica to leave (e). Rachel gets upset and ostensibly leaves the room (f).

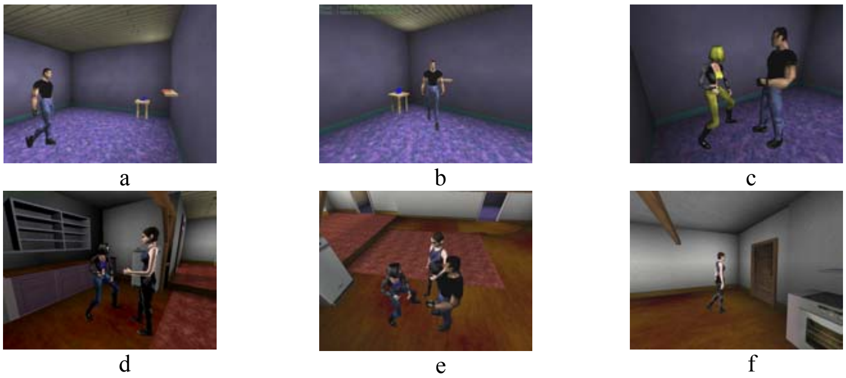

Fig 4. Sequence of actions illustrating the story instantiation

Let us now give a more technical description of these events, by detailing the associated steps in plan generation or terminal actions. Each of the main characters has its own planning system: they are synchronised through Unreal™ low-level mechanisms. Firstly, Ross' plan. The first sub-goal for Ross' plan is to acquire information about Rachel. There are various ways to satisfy this goal in Ross' behaviour representation, and the first one selected is to read her diary. The corresponding script involves going to the diary location and reading it. When Ross arrives in sight of the diary, the pre-condition of the action of "reading it" is checked: the diary is in place. This pre-condition is not satisfied, as the user intervened by removing the object from the set. Hence the second terminal action "ReadDiary" fails, as well as the whole sub-plan. The re-planning produces a new partial solution, which consists in asking Phoebe. Ross then goes to Phoebe's location and starts talking to her. As Phoebe is a reactive actor, she responds directly to Ross' request, in this case positively. In the meantime, Rachel's plan that governs her spontaneous activity, determines her to talk to her friend. She reaches Monica and starts conversing through a durative action (a scripted action which is associated a clock based on the internal Unreal™ clock). When Ross has finished talking to Phoebe, he needs to isolate Rachel in order to ask her out. The pre-conditions for a terminal action involving conversation with another actor is to check whether this actor is free. The personality profile defined initially for Ross (character with no ruthless manners) influences the heuristic values of the sub-plan nodes. So, Ross interrupts Rachel's conversation and, in a rude way, asks Monica to leave. Rachel reacts consequently to the situation by displaying a relevant mood state: she gets upset. Internally, the mood state is altered accordingly: all heuristics are revised, and of course, the activity "Talk to Monica" fails. Rachel leaves the room. In the same way, Ross' low-level mechanisms will provide situational information that will modify his internal states and influence his sub-plans. Ross will run after her when he realises Rachel is upset.

To summarise, this example illustrates the interaction of the two main character's plans, also influenced by user interference. Though these plans are designed from global narrative principles (considering the story genre), they are run independently. The particular interactions that take place depend on a number of variable factors, which contribute to the diversity of plots generated.

5 Conclusion

We have described an interactive storytelling system, which attempts to reconcile the character-based approach with the use of sophisticated narrative formalisms. At the heart of our system is a specific conception of user interaction. Namely, that the user can alter the ongoing events within the limits of the narrative genre itself. This amounts to saying that the entertaining aspects derive from some form of empowerment of the user, not to be passively linked to the plot: this form of interactivity is also a consequence of our emphasis on narrative structures.

There are many challenges related to this approach, in particular in further exploring narrative theories, with the prospect of implementing complex narratives and multiple storylines. The inclusion of more sophisticated language technologies in interactive storytelling is also a natural long-term goal, especially because of the relations between natural language semantics and narrative structures.

References

1. Dautenhahn, K.: Story-Telling in Virtual Environments, ECAI'98 Workshop on Intelligent Virtual Environments, Brighton, UK (1998).
2. Aylett, R.: Narrative in Virtual Environments: Towards Emergent Narrative. Papers from the AAAI Fall Symposium on Narrative Intelligence, Technical Report FS-99-01, AAAI Press (1999).
3. Sgouros, N.M., Papakonstantinou, G. and Tsanakas, P.: A Framework for Plot Control in Interactive Story Systems, Proceedings AAAI'96, Portland, AAAI Press (1996).
4. Young, R.M.: Creating Interactive Narrative Structures: The Potential for AI Approaches. AAAI Spring Symposium on Artificial Intelligence and Computer Games, AAAI Press (2000).
5. 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).
6. Nakatsu, R. and Tosa, N.: Interactive Movies, In: B. Furht (Ed), Handbook of Internet and Multimedia – Systems and applications, CRC Press and IEEE Press (1999).

7. 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).
8. 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).
9. Pirandello, L.: Sei personaggi in cerca d'autore (in Italian). Mondadori, Milan (1993).
10. Raskin, R: Five Parameters for Story Design. P.O.V., n. 5 (1998).
11. Cavazza, M. Charles, F., Mead, S.J., and Strachan, A.I.: Virtual Actors' Behaviour for 3D Interactive Storytelling. Proceedings of the Eurographics 2001 Conference (Short Paper), Manchester, UK, to appear (2001).
12. 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).
13. Webber, B., Badler, N., Di Eugenio, B., Geib, C., Levison, L. and Moore, M.: Instructions, Intentions and Expectations, Artificial Intelligence Journal, 73, pp. 253-269 (1995).
14. Prada, R., Machado, I., and Paiva, A.: Teatrix: Virtual Environments for Story Creation. Proceedings of Intelligent Tutoring Systems 2000, Canada (2000).
15. Barthes, R.: Introduction a l'Analyse Structurale des Récits (in French), Communications, 8, pp. 1-27 (1966).
16. Barthes, R.: S/Z (in French). Editions du Seuil, Paris (1970).
17. Laird, J.E.: It Knows What You're Going To Do: Adding Anticipation to a Quakebot, Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment, Technical Report SS-00-02, AAAI Press (2000).
18. Funge, J., Tu, X., and Terzopoulos, D.: Cognitive Modelling: Knowledge, Reasoning and Planning for Intelligent Characters. Proceedings of SIGGRAPH'99, Los Angeles, USA, pp. 29-38 (1999).
19. Tsuneto, R., Nau, D. and Hendler, J.: Plan-Refinement Strategies and Search-Space Size. Proceedings of the European Conference on Planning, pp. 414-426 (1997).
20. 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).
21. Geib, C.: The Intentional Planning System: ItPlans, Proceedings of the 2nd Artificial Intelligence Planning Conference (AIPS-94), pp. 55-64 (1994).
22. 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).
23. Charles, F., Mead, S. and Cavazza, M.: User Intervention in Virtual Interactive Storytelling. Proceedings of VRIC 2001, Laval, France (2001).