

Emergent Situations in Interactive Storytelling

Marc Cavazza, Fred Charles, Steven J. Mead
University of Teesside, School of Computing and Mathematics
Middlesbrough, TS1 3BA, United Kingdom
{m.o.cavazza, f.charles, steven.j.mead}@tees.ac.uk

Keywords

Interactive Storytelling, Virtual Humans, Autonomous Agents, Computer Games, Artificial Intelligence.

ABSTRACT

Interactive storytelling can either be based on explicit plot representations or on the autonomous behaviour of artificial characters. In such a character-based approach, the dynamic interaction between characters generates the actual plot from a generic storyline. Characters' behaviours are implemented through real-time search-based planning techniques. However, the top-down planning systems that control artificial actors need to be complemented with appropriate mechanisms dealing with emerging ("bottom-up") situations of narrative relevance. After discussing the determinants of plot variability and the mechanisms that account for the emergence of narrative situations, we introduce additional mechanisms for coping with these situations. These comprise situated reasoning and action repair: we most specifically illustrate the latter through a detailed example.

1. INTRODUCTION

In recent years, several paradigms have emerged for interactive storytelling [1]. They differ in both technical and artistic approaches. Within these, AI-based storytelling is more specifically concerned with the mechanisms for automatic story generation, which can be based on autonomous artificial actors' [2] [3] or on explicit plot representations [4]. These mechanisms support, in a unified framework, both story generation and user intervention.

We have been developing an interactive storytelling prototype based on autonomous characters. The specific application we are pursuing consists, for the spectator of a well-defined story (sitcom, drama, etc.), in being able to alter the plot while strictly preserving the story genre. Our system is based on the real-time generation of narratively meaningful situations featuring artificial actors, whose behaviours are produced by a real-time planning system, interleaving planning and execution. In this paper, we specifically discuss the various mechanisms underlying story generation that contribute to plot diversity. In storytelling systems

based on explicit plot representations [4] or narrative formalisations [5], story diversity can be directly derived from the representation of narrative functions. On the other hand, in character-based approaches, story diversity emerges from dynamic interaction between characters [6]. In other words, the narrative representation is based on roles rather than narrative functions [7]. Because of the great diversity of situations that can be generated by the interaction of artificial actors, there is a need to incorporate some kind of narrative control [3] [8] on top of the system, to ensure the narrative relevance of the generated situations. We have been exploring a complementary route, which consists in devising specific mechanisms to cope with the situations produced by the on-stage interactions between characters.

In the next sections, after a brief presentation of the system architecture, we discuss the main mechanisms for story generation and how these contribute to creating situations of narrative relevance. We then introduce specific techniques to cope with these situations in order to maintain narrative coherence. Finally, we illustrate these various aspects with results from our current prototype (Figure 1).

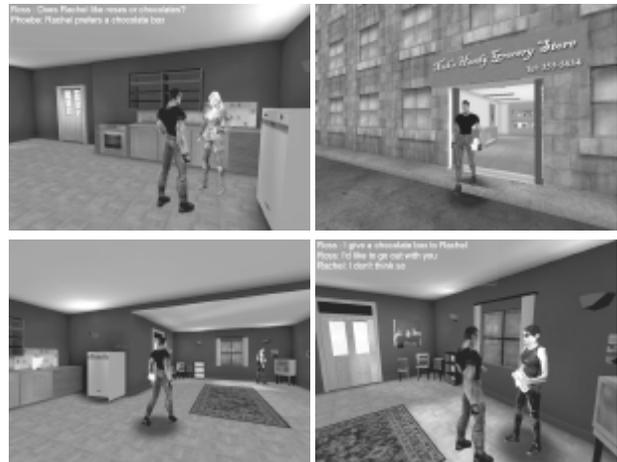


Figure 1. A story instantiation generated by the system.

2. INTERACTIVE STORYTELLING: SYSTEM DESCRIPTION

Our test scenario is inspired from a popular TV sitcom: in the small episode supporting our experiments, the main character "Ross" wants to invite the main female character "Rachel" on a date. The sitcom genre is an interesting test case in interactive storytelling, as both the ending and intervening situations are relevant. The system (Figure 2) is developed on top of the Unreal Tournament™ game engine: the set and characters have been designed (or imported) using the engine's development

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC2002, Madrid, Spain.

Copyright 2002 ACM 1-58113-445-2/02/03...\$5.00.

environment. All the AI modules supporting interactive storytelling have been developed in C++ (character’s behaviour) or UnrealScript™, the Unreal Tournament™ scripting language (interaction components).

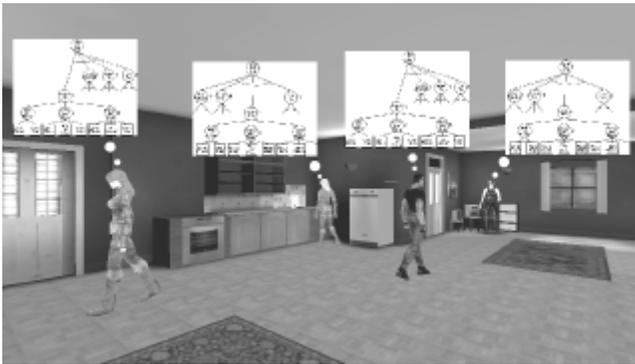


Figure 2. System architecture.

The system implements characters’ behaviours using AI planning techniques. In this context, Plans can be seen as a generic formalism for an actor’s behaviour and as a resource for story generation. They represent the storyline acted from a given character’s perspective, i.e. the character’s role. The basic hypothesis is that the overall story will emerge from the relations that exist between the various actors’ plans, these relations being determined from the story genre¹. For instance, if Ross’ plan is to seduce Rachel and Rachel’s plan just consists in carrying on her daily activities unaware of Ross, this is likely to result in a series of comic misunderstandings. We have hence defined separate plans for Ross and Rachel, which are in agreement with properties of the sitcom genre. Ross plan is to invite Rachel out for dinner. This plan is decomposed into a first set of high-level sub-goals: acquiring information about Rachel, attracting her attention, finding a way to talk to her privately, etc. On the other hand, Rachel’s plan is not specifically oriented towards Ross. Her plan will lead her to carry various activities, socially or privately, as a function of her mood and sociability.

We describe a character’s plan using a Hierarchical Task Network (Figure 3), which is formalised as an AND/OR graph. As we are representing narrative content a priori, our representations are actually *explicit* graphs (and this has implications for their automatic processing). From a formal perspective, the search process that is carried out by an AI planner takes an AND/OR graph and generates from it an equivalent state-space graph [9]. The process by which a state-space graph is normally produced from a Hierarchical Task Network (HTN) is called *serialisation* [10]. However, when the various sub-goals are independent from one another, the planner can build a solution straightforwardly by directly searching the AND/OR graph without the need for serialising it [10]. Further, there has been recently a renewed interest in search-based planning techniques, as these have demonstrated significant performance on various planning tasks [10] [11] [12] [13].

We thus use a real-time variant of the AO* algorithm [9] [14] [15] to search the AND/OR graph. The AO* algorithm is a heuristic

search algorithm operating on AND/OR graphs: it can find an optimal solution sub-graph, which in our case corresponds to a given character’s role in an instantiated story. The actual choice of sub-goal will depend on the heuristic value of each of these sub-goals, which contains narrative knowledge, such as the actor’s personality or “mood” (the latter value can be revised in a dynamic fashion).

In interactive storytelling, several actors, or the user himself, might interfere with one agent’s plans, causing its planned actions to fail. Hence, the story can only carry forward if the agent has re-planning capabilities. In any case, failed actions cannot be undone, as they have been played on stage. Action failure is indeed part of the story itself. This is why the dramatisation of actions must take their possible failure into account and store corresponding animations.

Considering the need for anytime interaction, we have developed a “real-time” variant of AO* that does not compute a complete solution sub-graph but interleaves planning and execution and only computes the partial solution tree required to carry out the next action. It explores the tree in a depth-first, left-to-right fashion [13] using essentially the heuristic part of evaluation functions. Story generation emerges from the interaction (“cross-product”) of the actors’ plans. While the story genre prescribes the overall relations between the main characters’ plans, there is no active synchronisation or prescribed dynamic interaction between these plans: their interaction only takes place through the events taking place in the virtual world.

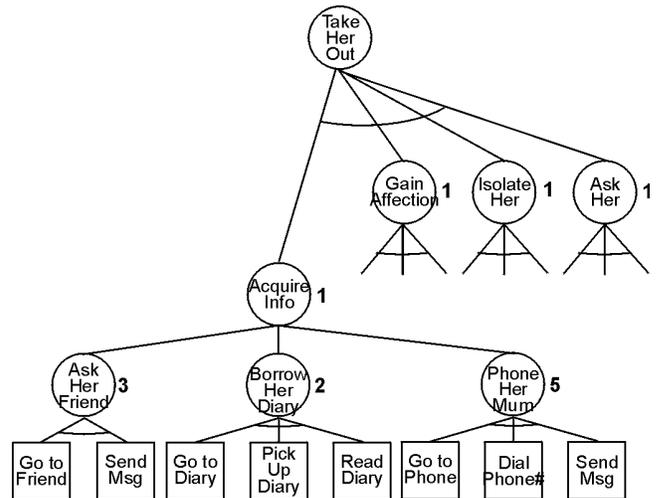


Figure 3. A character’s plan.

The interleaving of planning and execution also supports the global “interaction cycle” required for interactive storytelling. This cycle comprises: i) the dramatization of character’s actions, which convey their narrative meaning to the user (for instance, the fact that Ross reaches towards Rachel’s diary signals to the spectator the narrative importance of the object), ii) the potential intervention of the user in the story, either by acting upon on-stage objects, or by influencing virtual actors using speech recognition and iii) the staging of situations featuring interaction between actors. To understand the examples of this paper, a brief outline of Ross’ plan is necessary. In order to take Rachel out, Ross must first acquire information on her, such as availability and tastes. He should then find a way to talk to her in private,

¹ We could refer to this interaction as the “cross-product” of the individual characters’ plans.

which requires that he can reach her and that she be available. In the meantime, he should have gained her friendship or at least not alienated her (e.g. by upsetting her friends, etc.). He can offer her various gifts that will have to match her preferences. We will not detail here the various ways in which these sub-goals can be satisfied nor the executability conditions for the actions that form part of the sub-plans satisfying these goals, but this will become self-explanatory in the examples below.

3. DETERMINANTS OF NON-PREDICTABILITY IN STORY GENERATION

With the inherent tension between interactivity and storytelling [2] [8] [16], a balance has to be found between dramatisation (the narrative presentation), which should convey meaningful details to the user, and the generation of novel situations. For instance, it should be visible what kind of objects are relevant from a narrative perspective, such as keys, letters, flowers or a handgun (Footnote 2)². All these would be targets for user intervention. However, the precise outcome of user intervention should not be directly accessible to the user, in order to preserve the richness of story generation. In other words, stealing the handgun from the set should lead either to the character abandoning its murder plans or to it finding other means, depending on the circumstances. This, in our view, constitutes yet another principle of interactive storytelling: that interventions can have an impact at the plot development level (substantial alteration of the unfolding and the ending, such as murder vs. no murder) or at the situational level (the means by which goals are pursued and their associated dramatisation).

Narrative generation can thus be seen as comprising both a top-down and a bottom-up component, where the top-down part corresponds to the character's role and the bottom-up one consists in the situations created by interaction between characters, these not being determined a priori by their initial roles. In terms of story diversity, the top-down aspects are related to re-planning, i.e., choosing a new course of action when some attempt fails. Even though action failure can bear a dramatic interest, these aspects alone would not suffice to produce a wide range of narrative situations.

Conversely, the bottom-up aspects derive from the spatio-temporal relations between the various actors, which result in narratively meaningful situations. Examples of such situations are Monica standing in the room where the diary is, preventing Ross from stealing it, or Ross talking to Phoebe, which has the potential to make Rachel jealous. Even though the individual actors' behaviours, as supported through their underlying plans, are formally deterministic, there are several factors that contribute to the non-predictability of the plot at a macroscopic level, which is also the users' perspective. These are: i) the initial spatial allocation of the virtual actors, ii) the duration of actors' actions iii) the interaction between actors' plans, iv) the random outcome of some terminal actions and, naturally, v) user intervention.

In the next sections, we describe those factors that are the most relevant for the emergence of narrative situations.

3.1 INITIAL POSITION, ACTION DURATION AND SPEED OF ACTORS

At the beginning of the story, the characters are allocated random initial positions on the virtual set (they are "spawned" in Unreal™ terms). Their plans are then triggered from these initial positions: hence, they might direct themselves towards the on-stage resources required at the early stage of their plan. For instance, if Phoebe's first activity is to do some shopping, she would have to traverse the whole flat before reaching the exit. This leaves many opportunities for Ross to "intercept" her if necessary. On the other hand, if her initial position happens to be very close to the exit, she will leave the set in no time (Figure 4) and, if Ross needed to talk to her, (e.g., to get information about Phoebe) he would have to find another alternative (i.e., through replanning).

There is an obvious inter-dependence between initial spatial location and characters' intrinsic speed in carrying elementary actions and displacements. For instance, depending on their initial locations and respective speeds, Ross might or not be able to catch up with Phoebe before she leaves the flat for some shopping. This will in turn affect his obtaining information about Rachel, from which a different course of action will follow. At this stage, speed and duration of various actions is not a parameter the system can take into account, but that might become relevant in future developments.



Figure 4. Influence of spatial location.

3.2 USER INTERVENTION

User intervention is by nature non-deterministic and is thus a powerful factor of variability in story generation. It takes place as part of the overall interaction cycle: the user sees the action unfolding and, from the understanding he has gained from it, decides whether to interfere, and by which means. At this stage, the two modes of intervention are interaction through on-stage narrative objects and transmission of information to characters using speech recognition [18]. Intervention on objects will mostly result in action failure, namely failure of those narrative actions for which the object was a resource. A classical example consists in stealing Rachel's diary before Ross can get hold of it [7] (Figure 5). This will force Ross to re-plan a solution to the problem of acquiring information about Rachel. In the course of this new plan, many different situations can emerge that can trigger a "chain reaction" of causal events dramatically altering the course of action. While user intervention is possible at anytime, it is obviously constrained by the unfolding of the story itself. No retrospective intervention is allowed and the user can

2 In narrative studies, these are sometimes called "dispatchers", after Barthes' analysis of Tomachevski's narrative theory [17].

only interfere with the executability conditions of actions yet to take place. This has implications on how the user should follow the action when it is distributed through the virtual stage. Default cameras are centred on the main characters, such as Ross, but the user is also allowed to change viewpoints or to freely move in the environment.



Figure 5. Dramatisation of action failure due to user intervention.

3.3 INTERACTION BETWEEN CHARACTERS

Because characters' behaviours are determined by plans, the basis for characters' interaction is their competing for action resources, as part of the execution of their plans [6]. For instance, Ross might want to learn more about Rachel by reading her diary, but she might be using it herself, preventing Ross to do so (Figure 6). He could try to talk to Phoebe, but she'd be busy talking to Monica. In that sense, characters themselves can be seen as resources for other characters' actions: This competition for resources has the potential to trigger a causal "chain reaction" (though causality is not explicitly represented in a character-based approach). One classical example consists in competition for resources used in entertainment activities, such as TV sets or coffee machines, which have specific locations on stage and hence play an important role in the localisation of actors on stage. If a character is prevented from having access to some leisure equipment (because it is used by another actor, or has been moved by the user), it will have to re-plan another activity. In doing so, it will often move across the stage, which increases the probability of dynamic encounters with other characters, resulting in a whole range of situations. The latter point is supported by the fact that characters are aware of each other at various stages of the narrative action.



Figure 6. Competition for resources.

The existence of multiple actors naturally increases the probability of competition for resources and the generation of situations. In our current prototype, we have incorporated four autonomous actors, each with their own plan-based behaviour. Apart from competition for action resources, the interaction between characters' plans results in "random" on-stage encounters between agents that have the potential to create situations of narrative relevance. These constitute the "bottom-up" aspect of interactive storytelling: as it is not taken into account by plan-based behaviours, there is a need for specific mechanisms. These are situated reasoning and action repair.

4. SITUATED REASONING

The origins of situated reasoning in plan-based actors' behaviours [19] lie in the discrepancy between an agent's expectations and action preconditions.

One defining aspect of situated reasoning is that it is oriented towards obtaining a specific resulting state in a given situation [19]. We extend this definition by including that *avoiding an undesirable result should also be part of situated reasoning*. One such example in interactive storytelling consists in reacting to situations that emerge from the spatial interactions of artificial actors. As we have seen, the characters are allocated random positions on the set at the beginning of the story. As a consequence, while following their independent plans, this might result in situations that are not (and cannot be) explicitly represented as part of the plans, but cannot be ignored by the system.

One classical example consists in Ross meeting Rachel by accident while he is still at the early phase of his plan. (Figure 7). While he can choose to talk to her or to hide from her, one thing he definitely cannot do from a narrative perspective is to just walk past her without any interaction. One of the options offered by situated reasoning could be to hide from her, which can be implemented by interrupting Ross' current action. This also makes possible to naturally resume his initial plan: if Ross' current action was to meet Phoebe, he can return to her once Rachel has passed, not noticing him. In this specific case, hiding from Rachel does not impair sub-plan continuation. This might not be always the case, though. Let us consider a similar case, where Ross wants to talk to Phoebe but should not let Rachel see him talk to her, because he's afraid she could be jealous (a feature actually implemented in the system). He might wait, but unlike the

diary, Phoebe can in the meantime move to another location, or engage in other activities that would cause the initial intended action to fail. The interruption caused by situated reasoning can thus have an irreversible impact on the initial plan whenever there are time/duration or location constraints. However, even in this case, the use of situated reasoning has preserved the relevance and coherence of the plot, as situated reasoning (e.g., hiding from Rachel) is properly dramatised and constitutes a part of the story itself.



Figure 7. The emergence of situations.

One of the main causes for action failure is the non-satisfaction of executability conditions [19]. It is interesting to consider the case where the executability conditions are affected by the other agents' behaviours. One example consists in Ross needing to access Rachel's diary at the early stage of the story. This action can fail in several cases (corresponding to different contexts): the user has hidden the diary, Rachel herself is writing in the diary, or Monica is in the same room and hence Ross cannot steal it. The first case imposes re-planning, for action repair cannot be applied to the non-deterministic behaviour of the user (i.e., the user is unlikely to put the diary back). The second situation can be a target for action repair, as Ross could simply wait until Rachel has finished her task. More interestingly, the latter case offers the widest range of options. Ross can choose another source of information about Rachel, can wait for Monica to leave the room and resume his initial plan, or can try to influence Monica, so that he can still carry on his original action.

4.1 DIFFERENCE BETWEEN ACTION REPAIR AND SITUATED REASONING

There is sometimes a fine line between action repair and situated reasoning. Strictly speaking, action repair should be dedicated to recovering from action failure. However, in our storytelling context, action failure is most often due to the non-satisfaction of executability conditions due to external factors. For instance, Ross cannot read Rachel's diary because it is missing, because Rachel is using it, or because Monica is in the same room and he could not read it unnoticed. In other words, action repair is dedicated to restoring executability conditions and/or reaching the same final state as the original action, while situated reasoning essentially consists in interrupting the current plan and dealing with a specific situation arising. It hence does so more from the dramatisation perspective than from the planning perspective.

However, one of the fundamental questions is when to choose action repair over replanning. We can consider a specific

example. Ross wants to read Rachel's diary, but Phoebe is in the room. He thus cannot steal the diary (other variants are possible where Phoebe is sympathetic to this, but let us not consider these for the time being). He's got essentially two options: one is to look for another source of information about Rachel (re-planning), the other one is to repair the action by waiting for Phoebe to leave. Repair should be based on generic and principled knowledge about actions, such as the fact that presence has a limited duration in time (unlike absence). This problem also arises because re-planning in our storytelling context is essentially a short-range change of action focus, rather than a radically new strategy. Action repair can be used to restore executability conditions under certain circumstances, especially in the case of competition for action resources, which is typical of interaction between actors. We can consider the following example: because Rachel's diary has been stolen by the user, Ross needs to talk to Phoebe to acquire information about Rachel. However, Phoebe is engaged in a conversation with Monica in the shop. As he enters the shop, he notices that she is busy talking. He can restore her availability by either interrupting the conversation (which might upset Phoebe with unpredictable effects) or by waiting until she becomes available again. In the latter case, the waiting action is active and has to be dramatised. In other words, he will need to find some activity to perform until Phoebe becomes available. In the example of (Figure 8), he gets a drink from the vending machine and waits by it until Phoebe and Monica have finished talking. The dramatisation of events can also convey information about Ross' personality profile.

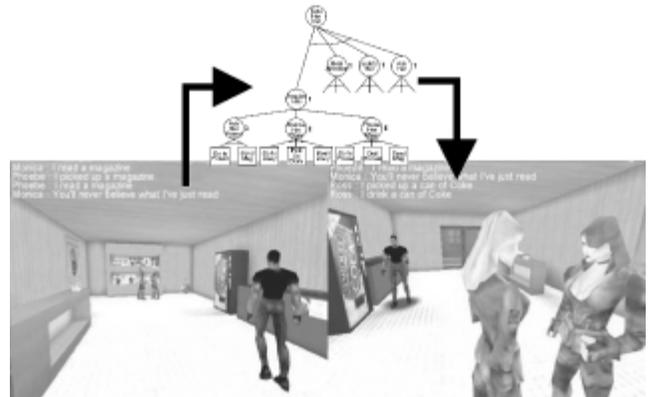


Figure 8. Action repair.

Action repair is implemented through a separate mechanism, which however shares the same formalism and techniques as standard behaviours. HTN graphs are used to describe the set of sub-plans, which will be used for repair by the character. The outcome of the sub-plan, developed as a consequence to the action repair, does not need to be explicitly transferred back to the characters' main plan. In action repair, unlike situation reasoning, there is no need to return post-conditions to the original plan when it resumes. The reason is that in most cases studied so far, repair was targeting executability conditions, which depended on narrative objects. Communication between the action repair modules and the generic plans takes place through the side effects of action resources, whether these be narrative objects or characters.

5. CONCLUSIONS

Character-based approaches in interactive storytelling have a good potential for story generation. Despite the deterministic nature of their underlying techniques, many different factors contribute to the non-predictability of the unfolding plot. In particular, while the user should be aware of the potential targets for interaction, the exact consequences of his intervention cannot be predicted. The top-down component of interactive storytelling, based on real-time planning for virtual characters, is however not sufficient to produce narratively meaningful stories. While some authors have proposed to introduce narrative control [3] [16] on the characters' or even the user's actions, there is also a need to cope with emerging situations that bear narrative relevance and cannot be incorporated into a priori representations. We have thus proposed to complement the top-down component of interactive storytelling with a bottom-up approach, in order to keep the original plans to a manageable size. This bottom-up component, following previous work in planning for embodied agents [19], comprises situated reasoning and action repair. Both can be implemented as standalone plans in the same formalism as the actors' original plans. Situated reasoning returns appropriate post-conditions when resuming the original plan, while action repair only operates on executability conditions and needs not updating any additional parameters. Further, action repair is an interesting alternative to re-planning and its practical implementation can incorporate several dramatic features that add value to the story presentation. These features should play an important role in the scaling-up of our current storytelling prototype.

6. REFERENCES

- [1] Mateas, M., and Sengers, P., 1999. Narrative Intelligence. AAAI Fall Symposium in Narrative Intelligence 1999, Cape Cod, MA, AAAI Press.
- [2] Young, R.M., 2000. Creating Interactive Narrative Structures: The Potential for AI Approaches. AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment 2000, AAAI Press.
- [3] Young, R.M., 2001. An Overview of the Mimesis Architecture: Integrating Narrative Control into a Gaming Environment. AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment 2001, AAAI Press.
- [4] 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.
- [5] Grabson and Braun, 2001 A Morphological Approach to Interactive Storytelling, Proceedings of Cast01, Living in Mixed Realities, Sankt Augustin, Germany, pp. 337-340.
- [6] Cavazza, M., Charles, F. and Mead, S.J., 2001a. Agents' Interaction in Virtual Storytelling. Intelligent Virtual Agents 2001 Conference, Madrid, Spain.
- [7] Cavazza, M., Charles, F., and Mead, S.J., 2001b. Characters in Search of an Author: AI-based Virtual Storytelling. First International Conference on Virtual Storytelling (ICVS 2001), Avignon, France.
- [8] Mateas, M., 2000. A Neo-Aristotelian Theory of Interactive Drama. AAAI Spring Symposium in Artificial Intelligence and Interactive Entertainment 2000, AAAI Press.
- [9] Nilsson, N.J., 1998. Artificial Intelligence: A New Synthesis. San Francisco, Morgan Kaufmann.
- [10] Tsuneto, R., Nau, D. and Hendler, J., 1997. Plan-Refinement Strategies and Search-Space Size. Proceedings of the European Conference on Planning, pp. 414-426.
- [11] Bonet, B. and Geffner, H., 1999. Planning as Heuristic Search: New Results. Proceedings of ECP'99, pp. 360-372.
- [12] Korf, R.E., 1990. Real-time heuristic search. Artificial Intelligence, 42:2-3, pp. 189-211.
- [13] Pemberton, J.C. and Korf, R.E., 1994. Incremental Search Algorithms for Real-Time Decision Making. Proceedings of the 2nd Artificial Intelligence Planning Systems Conference (AIPS-94).
- [14] Pearl, J., 1984. Heuristics: Intelligent Search Strategies for Computer Problem Solving. Reading (Massachusetts), Addison-Wesley, 1984.
- [15] Knight, K. and Rich, E., 1991. Artificial Intelligence, 2nd Edition. McGraw Hill.
- [16] Mateas, M., 1997. 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.
- [17] Barthes, R. 1966. Introduction à l'Analyse Structurale des Récits (in French), Communications, 8, pp. 1-27.
- [18] Charles, F., Mead, S. and Cavazza, M., 2001. User Intervention in Virtual Interactive Storytelling. Proceedings of VRIC 2001, Laval, France.
- [19] Geib, C. and Webber, B., 1993 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.