Automated Control of Interactions in Virtual Spaces: a Useful Task for Exploratory Creativity

Federico Peinado, Miguel Ancochea and Pablo Gervás

Dep. Sistemas Informáticos y Programación
Universidad Complutense de Madrid, Spain
fpeinado@fdi.ucm.es mancochea@indra.es pgervas@sip.ucm.es
http:\\gaia.sip.ucm.es

Abstract. The main problem for Interactive Digital Storytelling is the inevitable conflict between author's determinism and interactor's freedom. This paper presents a computational solution for addressing this problem, and discusses how exploratory creativity (more than transformational creativity) can play a role in solving the dilemma. Our proposal is based on the theoretical study of tabletop Role-Playing Games, Virtual Environments applications and structural analysis of narrative texts. It involves the implementation of automatic directors for the *story* and *discourse* levels useful for many story-based applications.

1 Introduction

The use of physical metaphors to help with the organization of information presents important advantages from the point of view of navigability, ease of access to particular items of information, and orientation of the user while exploring a particular set of data. As more and more institutions become virtual (retail outlets for products such as books or computers, museums, universities, banks, libraries, newspapers...) this type of solution for presenting information is becoming increasingly popular. Even if it is not made explicit, a hypertext web site can be understood as a virtual space if particular documents are interpreted as rooms and hyperlinks as connections between them.

Designing a virtual space involves finding a trade-off between two possibly divergent aims: to include enough similarity to an existing concept to act as guidance for the user, and to exploit the advantages of the virtual space over the concepts which it mirrors (every available object can be as close as the next room, experiences can be explicitly tailored to each individual, displays can be dynamically rearranged based on user feedback...). In many cases, designers of web sites tend to focus on exploiting the advantages of the virtual space. This leads to web sites where the user can easily become lost, overwhelmed by the amount of information, or unable to find something that he knows must be somewhere close. As a result, several research initiatives have addressed the problems of guiding the user around this type of virtual space, particularly centered around electronic catalogues and virtual museums [7,15,18,21].

An extreme case of this type of situation is provided by the emerging field of Interactive Digital Storytelling, in which the concept of virtual space is blended with literary fiction. The main problem of Interactive Digital Storytelling (IDS) design is structuring an intensive and meaningful interactive experience at the same time as enabling the development of a good pre-authored plot. This is a real dilemma, because while the interactor of some story-based application is taking a lot of decisions about the performance of his character that change the plot development, the author's plan may have been developed according to a different storyline.

This paper presents a computational solution for addressing this problem, and discusses how exploratory creativity (more than transformational creativity) can play a role in solving the dilemma.

2 Relevant Techniques and Useful Sources

Many approaches are found in the literature that try to solve this conflict in an automated or semi-automated way. Basically, they make interactive storylines by adapting author's plot units or other pre-established resources to the interactor behaviour at run-time. This task requires a computational solution that can react appropriately to "unexpected" user decisions.

In order to do this, several aspects have to be taken into account: an adequate representation of the needs or preferences of particular users, a valid heuristic for finding new solutions based on existing ones, appropriate technology for rendering the messages to the user as natural language text, and some guidance as to how the control of the interaction may be bettered achieved. This section covers the technologies and sources used in this work to cover those aspects, as well as presenting an existing Interactive Fiction engine that is used as case study platform to test the solutions presented here.

2.1 User Modeling

In order to apply the personalization described above, it is crucial to identify correctly the character profile of the interactor. In a RPG, this task must be carried out based on limited information like the description of the character explicitly required by the rules, observation of the player's reactions, and possibly the player's record in previous games. To simulate this dynamically in an interactive system is the next step of this study. For our current purposes, it is enough to provide interactors with a set of predefined characters, such that each one of them is related with a specific interactor model. It is hoped that interactors of a specific model will under such conditions choose the type of character most related to their preferences in acting. The initial adscription of interactor model to character type will be used by the system to assign interactor models to the interactors.

Interactor Models for Story-Based Games Laws identifies seven basic types of role-players according to their motivation and the sort of characteristics that they expect of a game in order to consider satisfactory. These motivation characteristics are referred to as *plot hooks*.

Power Gamer searches for new abilities and special equipment.
Butt-Kicker always waits for a chance to fight with someone else.
Tactician feels happy in the anticlimax scenes, thinking about logical obstacles to overcome.
Specialist just wants to do the things that his favorite character do.
Method Actor needs situations that test his personality traits.
Storyteller looks for complex plot threads and action movement.
Casual Gamer remains in the background (and has a very low degree of participation in general).

Interactor Models for Virtual Museums It is widely accepted and many authors agree [21,18,15,17] that moving on a virtual space or environment as a virtual multimedia-based museum needs user modeling in order to provide the most convenient adaptation mechanisms to make every visitor enjoy a so-called "museum experience".

Keeping the visitor interested during all the route is essential, as it is also giving him/her the most suitable information, so that it is not repeated nor more than needed and coherent with the previously given information.

Aiming to model the museum visitor as closely to the real thing as possible there have been a series of initiatives in many fields to determine what are the main behavioural features to be considered of a person in that circumstances, just as it has been modeled on Role-Playing Games. In the HIPS approach [21], the authors take into a count a very useful classification based on previous psycho-social studies.

Ant visitor tends to follow the path proposed by the curator.
Fish visitor prefers to move in the centre of each room, without looking at details of the art, but cruising all the exhibition.
Butterfly visitor frequently changes direction, and don't follow the purposed path, although they manage to see almost all the artworks.
Grasshopper visitor sees only the pieces of artwork they are interested in, leaded by their personal interests and knowledge.

2.2 Knowledge-Intensitive Case-Based Reasoning

Knowledge Intensive Case-based reasoning relies on applying additional explicit knowledge to improve the performance of case-based reasonings approaches that rely mostly in reusing existing solutions to adapt to new problems. The COLIBRI (*Cases and Ontology Libraries Integration for Building Reasoning Infrastructures*) system assists during the design of KI-CBR systems that combine cases with various knowledge types and reasoning methods. It is based on CBROnto [8,9,10], an ontology that incorporates reusable CBR knowledge and serves as a domain-independent framework to develop CBR systems based on generic components like domain ontologies and Problem Solving Methods (PSMs).

2.3 Natural Language Generation Architecture

FROGS is a flexible object-oriented Java-based framework to build Natural Language Generation (NLG) applications taking RAGS [3] as a reference and thus implementing its main standard definitions for the abstract data model and using XML for the real source data.

This tool, which also provides a sample default implementation called *jGolen*, supports a wide selection of the most common state-of-the-art generation architectures, ranging from a simple monolithic implementation to a revision-based architecture or a blackboard and including the frequently used Reiter's pipeline [23] or the interactive feedback-based architecture.

These architectures usually develop on a series of common generation phases also provided by the framework, including content determination, discourse planning, sentence aggregation, lexicalization or linguistic realization.

2.4 Basic Source for Control Rules

Because Interactive Narrative is a relatively new field and it is difficult to find formal studies about interactive plot development, we have done a review over the modern methodology of Role-Playing Games (RPGs), the interaction in Virtual Environments (VEs) and the classic studies about narrative structuralism like *Story and Discourse* by Seymour Chatman [5].

Tabletop RPGs are exercises in intellect and imagination: a group of players sitting around a table, rolling dices and playing out an imaginary role in a complex shared fantasy, true collaborative narrative.

The Game Master (GM) is a special kind of player, he is the "interactive storyteller". He designs all the elements of the story and he manage all the possible events that can occur in its development, improvising the dialogue contributions of non-player characters, resolving players actions, etc.

The degree of interactivity in RPG can be enormous, limited only by the players imagination. This implies that no GM, however experienced, can have a deep enough plan to react appropriately to all the possible actions that players can come up with in the world of fiction. To operate successfully without such a plan, GMs must use their imagination, improvise adequate solutions, and continuously rewrite their plots on the fly.

The figure of GM is the best model we have found in real life for designing and directing interactive stories. For the development of the work presented here we have used the description of the relevant heuristics given by Robin Laws [13].

2.5 Internet Adventure Game Engine: A Case Study

Text Adventure Games, broadly known as the Interactive Fiction genre, appeared as the first narrative games at the end of 70's. Originally, interactive fictions are like interactive books, only made of text chunks. They have complex plots and offers a narrative presentation to the player. In this kind of applications, story and world simulation are tightly coupled.

Internet Adventure Game Engine (IAGE [22]) is a Java Open Source project for the development of a multiplayer interactive fiction system. In contrast to Massive Multiplayer Online Role-Playing Games (MMORPGs), which maintain a lot of players playing at the same time, with as many ongoing stories as users connected to the server, IAGE allows one pre-authored storyline with the added possibility of having more than one player immersed in the same story. IAGE can be also used to create single player applications like traditional systems as Inform [19].

3 Story and Discourse Architecture

Automated control of interactive narrative requires a system which is able to improvise acceptable and engaging solutions to conflicts arising at two different levels. One is the creation of characters, places or situations that will appear in the *story* and the other is related to the form of the messages that the interactor will send and receive during the *discourse* generation in real time (specially in case of natural language interfaces). At each level, the definition of what is considered acceptable and engaging will vary.

Our approach is based on developing a Case-Based Reasoning (CBR) model of a particular set of algorithms and heuristics for RPG mastering and VE interaction, and applying this model to narrate the experience of each interactor in a multiuser & directed i-fiction engine.

We propose a global architecture for a system which employs both CBR-driven story management, user modelling and adaptative context-sensitive discourse generation.

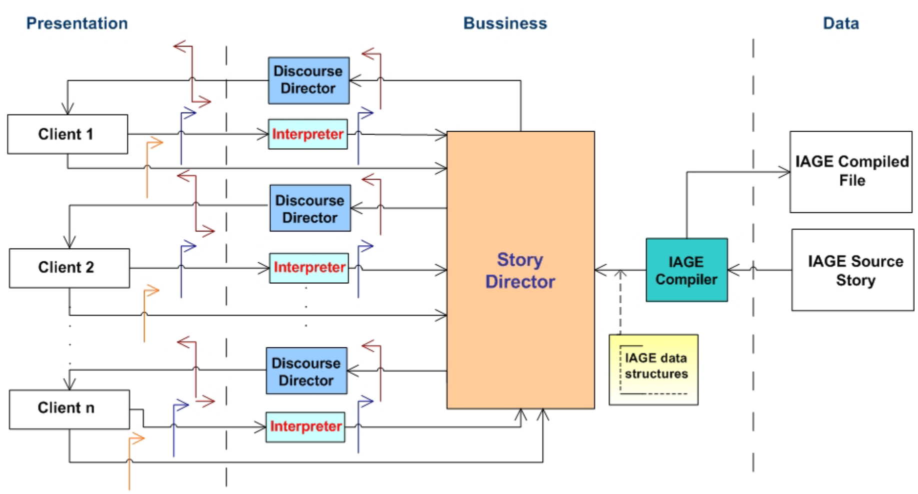

Fig. 1. Overview of the multitier client/server IDS system architecture

Firstly, there would be a *story director* as a standalone component. Like an agent *agora*, this module is intended to provide adaptative management of the VE and interactions for every "agent" on the systen, including the characters and the *discourse director* components.

This last component is created once for every interactor and attached to his client application by the story director. Its aim is not only to act as an NLG module, but also adapt generation to the interactor model and to the current state of the VE just as the interactor percieves it.

The story director performs a previous selection of the environment currently available for each interactor, as a graph of facts and entities, and gives it to discourse directors, enriched with the history of previous interactions. Discourse directors should then perform language generations which try to attract and interest the interactor, at least as the system can see him.

There would also be modules in charge of user input processing or interpreters, with raw natural language processing capabilities.

As a technological standard of reference to build a system based on this architecture, we propose that of a J2EE multitier application. This way, there would be a lightweight-client presentation layer, a bussiness server layer involving all story an discourse management processes and a back-end server data layer, as shown on Figure 1.

Interfaces are shown as arrows pointing to the direction of the necessary implementor. Upon execution of a lightweight client application, it registers itself on the story director, which in turn instantiates a discourse director and an interpreter and attaches them to the client application, so that it only communicates with them from then on. Each interpreter processes input, passes it to the story director and it decides whether to notify to the corresponding discourse director any relevant change on the VE or interactor, thus implementing an architectural Observer pattern, on which the story director is the subject to be observed. In order to unbound this communication, an additional Factory pattern should be employed to instantiate discourse directors for the story one.

3.1 Case-Based Reasoning Story Director

In searching for the right computational solution, widely different approaches have been proposed by researchers in the area. For example, there are proposals based on dynamic behavior of autonomous characters that achieve dramatic goals [4,25]. Other approaches give more responsibility to a central dramatic planning algorithm, using directable characters [14,11] or just a stand-alone dramatic planner that controls the most important narrative elements, like characters [16] or the whole fictional world [12]. In the work of [11] the CBR full life cycle -retrieval, adaptation, reuse and repair of previous solutions for new problems- is used for storyline representation and a strategy formalization that allows for storyline adaptation.

We propose a Knowledge Intensive CBR (KI-CBR) approach to the problem of generating interactive stories from the game state and a set of cases specifying how the next scene will be performed. In our model of Interactive Storytelling, adapted from the original RPG conventions, we separate the world simulation from the story control. The IDS system that we are considering has a narrative environment with a high level of interactivity (textual commands) and uses IAGE as a low level world simulator. Over that we have a CBR system that can guide the development of plot, building creative new situations on the fly from the case base.

Knowlege Representation In order to be able to use the Chatman's concepts and the user models computationally, they have been translated into an ontology that gives semantic coherence and structure to our case base.

An additional ontology provides the background knowledge required by the system, as well as the respective information about characters, places and objects of our world. This is used to measure the semantical distance between similar cases or situations, and maintaining a independent story structure from the simulated world.

The author fills the case base with scenes that contain preconditions and postconditions. Because the case base is made using cases proposed by the author, we know that the system makes variants of many elements that can be combined in a lot of different ways.

Plot Driving Algorithm Our plot driving algorithm is based on a CBR interpretation of the Law's improvising method [13]. This method is based on making transitions to the most *interesting* story scenes whenever it is possible.

The CBR system uses two similarity functions. The first one is used to recover the scene that leads to the most obvious/surprising transition from the current scene. Every case has a obvious/surprising property and the author assigns negative values to obvious scenes and positive values to surprising ones.

The second similarity function is used to retrieve the scene that involves the most pleasing/challenging transition from the current scene. The definition of *how pleasing/challenging a scene is* is given by the number of easy/difficult tasks matching the interactor plot hooks that appears in the scene. The author selects one kind of interactor as a reference and assigns negative values to pleasing scenes and positive values to challenging ones, based on those things that the interactor prefers.

For example, an ambush during a walk in the Palace's garden is a surprising scene, a big challenge for a Butt-Kicker playing an adventure game. In the other hand, free exploration of the museum shop at the end of the visit is an obvious but pleasing transition for a Butterfly visitor.

In this way the algorithm includes a number of obvious paths and other paths that may progressively get more positive for the interactor interests. Additionally, it may include surprising or negative scenes.

3.2 Natural Language Discourse Director

On a general, architectural basis, it is a relatively complex module of adaptative generation. Its core is intended to be standalone and not bound to a domain or the global system on which it is used, making therefore necessary a domain adapter or wrapper. This adapter processes its domain-dependent input and builds a series of generation and user model hints to be passed to the core component.

The core itself consists of a component of raw language generation and a module to adapt the previous hints to the current user an environment context, a contextual adapter. Eventually, this new adapter will start generation with a series of hints and provide a textual planning algorithm by building a text-potential graph based on entities, facts and relations, similar to that found on ILEX [17], so that the core NLG plans easily by implementing a traverse algorithm for the graph.

Finally, the NLG core is intended to be an implementation of the FROGS framework, thus implementing also the standards proposed in RAGS as a reference architecture [3,2].

4 Discussion

In the present context, when we are considering the automation of the task of dynamically controlling the flow of interactive narrative, it becomes imperative to discuss at least two basic issues:

– whether this task (or parts of it) indeed involves creativity at some level
– how this particular type of creativity can be classified in terms of Boden's analysis [1]
– whether this kind of creativity can be automated in any way in a system of the kind described

It is clear that controlling interactive narrative, such as acting as Game Master in a Role Playing Game, is generally perceived as a creative task. Human candidates to carry out the task have to fulfil certain expectations on the part of the players, and knowledgeable players are quick to judge whether a game master is good or not. However, there is no easy way of extrapolating explicit rules as to how this evaluation is carried out.

On the other hand, the generation of literary fiction, even of the non-interactive variety, is considered creative. Art and music require the generation of artefacts that are radically creative, in the sense that they cannot be classified under existing genres. In these fields, an artefact is deemed creative only if it defines a style of its own. In the field of literature, the situation is more complex. The definition of a new genre is not an immediate aim of the community of creators. Rather, they tend to focus on being creative within given genres.

One possible explanation of this difference is related to the role of meaning in each of these fields. Music, and non figurative art, produce artefacts that have no intrinsic meaning. In literature, on the other hand, meaning is crucial, meaning is probably a very high percentage of the value of the artefact. This is not the case in certain types of free style poetry.

Another possible factor to take into account is utility. Wherever artefacts are expected to be useful, or to fulfil a specific set of needs, the kind (or the degree) of creativity expected is reduced. In a way, a particular set of needs to be fulfilled constitutes a skeleton definition of a genre.

This has become particularly apparent with the mass production and dissemination of musical and literary material. People have come to expect certain artefacts to fill certain requirements. For instance, airport novels must grip the reader and keep him entertained over a reasonably long period while requiring little effort. Pop songs intended to be played in radio stations are expected to have a certain length, to be catchy, and to have refrains that get repeated a few times. In either case, there is an industry that sets the requirements and drives the type of material that is produced. Should this be taken to imply that no creativity is involved? This is not true. In fact, within each field, success depends very much in being creative within the given requirements, as this becomes the distinguishing feature of a particular product with respect to its competitors.

However, it could be interpreted to mean that the kind of creativity required in these cases is not transformational creativity, but rather exploratory creativity. In general, it is easy to assume that the amount of creativity required to produce valid artefacts in these fields is very low. We have all heard someone complain about the injustice of bestselling authors who write books by the rules and make millions out of them, having exercised very little "creativity". If this were true, it ought to be easy to program computers to carry out this kind of tasks. We have yet to hear about major breakthroughs of computer authored artefacts in any of these industries.

From the point of view of the feasibility of automating in any way the creativity involved in these tasks, it is important to take into account that the kind of creativity that is being sought is not transformational creativity but exploratory creativity. In the particular domain of interactive fiction, and narrowing down to the proposed case study, the amount of creativity required would be restricted to finding adequate combinations of whatever resources (locations, characters, objects, events, character functions...) are explicitly represented in the available ontology.

The results obtained by the system when carrying out the task of controlling real interactions can be evaluated in two ways. On one hand, they can be compared with the results of equivalente tasks carried out by human operators. To this end it is important to provide the system with adequate facilities for keeping logs of particular games. This evaluation needs the human users intervention. On the other hand, attempts can be made to evaluate from a more formal perspective the level of creativity actually displayed by the system. This evaluation can be oriented towards locating any indications of creativity introduced by the process in comparison with the original samples which make up the case base, possibly by applying metrics and analyses of creative activities that have progressively emerged over the recent years [24,6,20].

5 Conclusions

The task described in this paper constitutes a good example of a circumstance where exploratory creativity appears not as a "lesser sister" of transformational creativity but as a more adequate alternative. The actual requirements imposed by the task suggest that a solution based on exploratory creativity rather than transformational creativity would be more adequate, and that the level of creativity required may be susceptible of automation to a certain degree. As in many cases where creativity is required, some degree of aesthetical value is needed for the results obtained to be considered acceptable. However, in the context of interactive narrative, the level of aesthetical quality from a literary point of view is quite low when compared to narrative in general fiction. Consumers of this type of narrative tend to focus more on issues like the sequence of events or the pace of the interaction than the aesthetic quality of system messages from a literary point of view. This presents many advantages since even results of low aestehtic quality (as may be obtained in the erly stages of teh development of the system) may still be valuable from the point of view of interactive narrative.

Although the system is not fully-implemented yet, the progress so far points to a reasonable solution for interactive narrative generation. As outlined in the introduction, once the system is fully developed, the approach presented here for interactive narrative may be extended to other situations where automated control of interactions in a virtual space is required. The design of the system has taken this possibility into account, by ensuring that adequate modularity is included in the design. In this way, domain specific resources such as user models are isolated in particular points of the code to allow for easy interchangeability.

As possible future applications of the resulting system techniques, we are considering two areas that involve interation in virtual spaces for which specific developments are currently under way in our department. One possibility is a Virtual Museum of Computer Science that is under construction in our university. This project involves the creation of a large virtual space in which various materials concerning computer science will be displayed. Our research group is keeping in contact with the museum's development team with a view towards ensuring compatibility in data formats. The aim is to develop an automatic museum guide that would monitor the user's navigation through the museum, providing customized generated comments and dynamically redesigning the floor plan to suit user needs. The other possibility concerns JV2M, a knowledge-based learning environments environment where students can learn the Java Virtual Machine (JVM) structure and Java language compilation. The system presents a metaphorical 3D virtual environment which simulates the JVM and the user is symbolized as an avatar which is used to interact with the virtual objects. An animated pedagogical agent called Javy (JavA taught VirtuallY) monitors the student whilst she is solving a problem with the purpose of detecting the errors she makes in order to give her advice or guidance. The current version of the JV2M system would clearly be enhanced by the application of the techniques considered here.

6 Acknowledgements

The first author is supported by a FPI Predoctoral Grant from Complutense University of Madrid. The work was partially funded by the Spanish Committee of Science & Technology (TIC2002-01961).

References

1. M. A. Boden. *The Creative Mind: Myths and Mechanisms*. Wedenfield and Nicholson Ltd, London, 1990.
2. L. Cahill, J. Carroll, R. Evans, D. Paiva, R. Power, D. Scott, and K. Van Deemter. From rags to riches: Exploiting the potential of a flexible generation architecture. In *Association for Computational Linguistics*, Toulouse, France, 2001.
3. L. Cahill, C. Doran, R. Evans, C. Mellish, D. Paiva, M. Reape, D. Scott, and T. N. In search of a reference architecture for nlg systems. In *7th European Workshop on Natural Language Generation*, pages 77–85, 1999.
4. M. Cavazza, F. Charles, and S. J. Mead. Character based interactive storytelling. *IEEE Intelligent Systems*, 17:17–22, 2002.
5. S. Chatman. *Story and Discourse : Narrative Structure in Fiction and Film*. Cornell University Press, second edition, 1986.
6. S. Colton, A. Pease, and G. Ritchie. The effect of input knowledge on creativity. In A. Cardoso, C. Bento, and G. Wiggins, editors, *First Workshop on Creative Systems, International Conference of Case-Based Reasoning*, 2001.
7. R. Dale, J. Oberlander, M. Milosavljevic, and A. Knott. Integrating natural language generation and hypertext to produce dynamic documents. *Interacting with Computers*, 11(2):109–135, 1998.
8. B. Díaz-Agudo and P. A. González-Calero. An architecture for knowledge intensive CBR systems. In E. Blanzieri and L. Portinale, editors, *Advances in Case-Based Reasoning*, New York, 2000. Springer-Verlag Berlin Heidelberg.
9. B. Díaz-Agudo and P. A. González Calero. A declarative similarity framework for knowledge intensive CBR. In *International Conference on Case-Based Reasoning*. Springer-Verlag, 2001.
10. B. Díaz-Agudo and P. A. González Calero. Knowledge intensive CBR through ontologies. *Expert Update*, 6(1), 2003.
11. C. Fairclough and P. Cunningham. A multiplayer case based story engine. In *4th International Conference on Intelligent Games and Simulation*, pages 41–46. EUROSIS, 2003.
12. A. S. Gordon and N. V. Iuppa. Experience management using storyline adaptation strategies. In *First International Conference on Technologies for Digital Storytelling and Entertainment*, Darmstadt, Germany, 2003.
13. R. D. Laws. *Robin's Laws of Good Game Mastering*. Steve Jackson Games, first edition, 2002.
14. S. C. Marsella, W. L. Johnson, and L. Catherine. Interactive pedagogical drama. In C. Sierra, M. Gini, and J. S. Rosenschein, editors, *Fourth International Conference on Autonomous Agents*, pages 301–308, Barcelona, Spain, 2000. ACM Press.
15. P. Marti, A. Rizzo, L. Petroni, G. Tozzi, and M. DiligentiI. Adapting the museum: A non-intrusive user modeling approach. In *User Modelling*, 1999.
16. M. Mateas and A. Stern. Façade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference, Game Design track*, 2003.

17. C. Mellish, M. ODonnell, J. Oberlander, and A. Knott. An architecture for opportunistic text generation. In *International Workshop on Natural Language Generation*, Niagra, 1998.
18. M. Milosavljevic, R. Dale, S. J. Green, C. Paris, and W. S. Virtual museums on the information superhighway: Prospects and potholes. In *Annual Conference of the International Committee for Documentation of the International Council of Museums*, 1998.
19. G. Nelson. Inform manual, 2001. http://www.inform-fiction.org/.
20. A. Pease, D. Winterstein, and S. Colton. Evaluating machine creativity. In A. Cardoso, C. Bento, and G. Wiggins, editors, *First Workshop on Creative Systems, International Conference of Case-Based Reasoning*, 2001.
21. D. Petrelli, E. Not, and M. Zancanaro. Getting engaged and getting tired: What is in a museum experience. In *User Modelling: Workshop on Attitude, Personality and Emotions in User-Adapted Interaction*, Banff, 1999.
22. R. Rawson-Tetley. Internet adventure game engine (IAGE), 2002. http://www.ifarchive.org/if-archive/programming/iage/.
23. E. Reiter. Has a consensus nl generation architecture appeared, and is it psycholinguistically plausible? In *Seventh International Workshop on Natural Language Generation*, pages 163–170, Kennebunkport, Maine, USA, 1994.
24. G. D. Ritchie. Assessing creativity. In G. Wiggins, editor, *Symposium on Artificial Intelligence and Creativity in Arts and Sciences*, 2001.
25. R. M. Young. Notes on the use of plan structure in the creation of interactive plot. Technical Report FS-99-01, AAAI Fall Symposium on Narrative Intelligence, AAAI Press, Menlo Park, 1999.