A Multi-User Natural Language Interface for Interactive Storytelling in TV and Cinema

Edirlei Soares de Lima [1] Bruno Feijó [1] Cesar T. Pozzer [2] Antonio L. Furtado [1]
Angelo E. M. Ciarlini [3] Fabio A. Guilherme da Silva [1]

[1]PUC-Rio, Departamento de Informática, Brasil
[2]UFSM, Departamento de Eletrônica e Computação, Brasil
[3]UNIRIO, Departamento de Informática Aplicada, Brasil

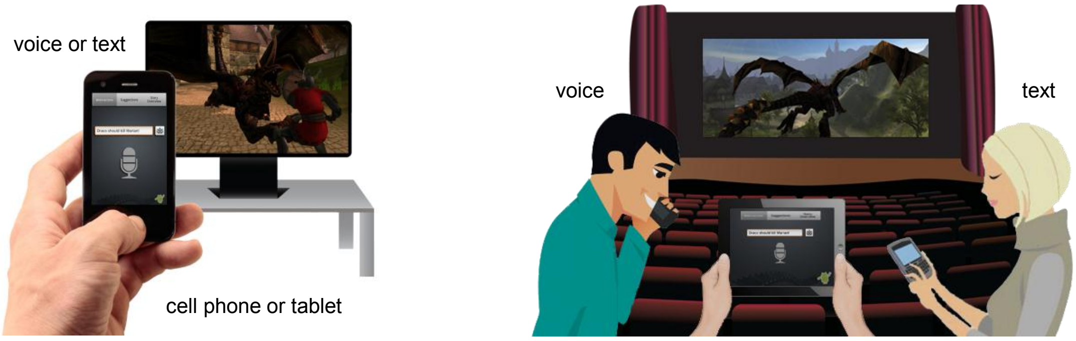

(a) Interactive TV (b) Interactive Cinema

Figure 1: The multi-user natural language interaction system for interactive storytelling.

Abstract

Interactive narratives designed for television and cinema demands new interaction mechanisms. In the TV context, thousands of viewers sharing the same narrative at different places must have the opportunity to interact with ongoing stories. The same requisite is valid for cinema, where the whole audience in the theater must have the opportunity to interact. This paper presents a multi-user natural language interface for interactive storytelling applications designed for TV and cinema. The proposed system combines the use of mobile devices with natural language to allow users to freely interact with virtual characters by text or speech.

Keywords: Interactive Storytelling, User Interface, Interactive TV, Interactive Cinema, Natural Language.

Authors' contact:
`{elima,bfeijo,furtado,faraujo}@inf.puc-rio.br,`
`pozzer@inf.ufsm.br,`
`angelo.ciarlini@uniriotec.br`

1. Introduction

The first interactive narratives can be dated back to the 1970s [Meehan 1977] and important experiments on agent-based storytelling systems took place in early 1990s [Loyall and Bates 1991]. In the 2000s we can find the most influential research works on interactive storytelling systems [Cavazza *et al.* 2002][Mateas and Stern 2003]. In more recent years, we have been exposed to new demands for richer interactive experiences in storytelling, such as transmedia storytelling [Cheshire and Burton 2010], social interaction between groups [Williams *et al.* 2011], and interactive TV [Ursu *et al.* 2008a]. These new interactive experiences require the development of especial interaction mechanisms.

An interaction mechanism defines how users are able to communicate their desires and expressing satisfaction with the ongoing story. These mechanisms should consider the characteristics of the media platform that supports the story. In television, for instance, interaction systems should consider one or few local viewers (in the same room) or thousands of viewers sharing the same story at different places. In movies, the audience is restricted to the theater space (although synchronous forms of projection could connect several theaters). Interactive narratives designed either for TV or cinema require new interaction systems that support multi-user interactions. Moreover, the interaction must be simple, comfortable and attractive to users. Ideally, the interaction system should be transparent, that is, it should not occupy the screen space with 2D widgets or any other pictorial artifacts that interfere with the visual language of the film or the TV program.

In this paper we propose a natural language interface for interactive TV/Cinema storytelling that accomplishes the requirements for a transparent multi-user interaction interface. It combines the use of mobile devices with natural language to allow users to freely interact with virtual characters by text or speech (Figure 1). We implemented a prototype of the proposed interaction system in Logtell [2012], an interactive storytelling system based on temporal logic and non-deterministic planning. The interactive story implemented in our system corresponds to a short story in the genre of swords and dragons using 3D characters (Figure 2). The dramatization module of Logtell can either work with 3D models or living actors [Lima 2012].

Figure 2: 3D dramatization of the swords-and-dragons story performed by the proposed system.

2. Previous Works

There are several works on forms of interactions for interactive storytelling in the literature. They cover the subject from traditional GUI interfaces [Ciarlini *et al.* 2005][Grasbon and Braun 2001] to more complex interaction mechanisms, such as speech recognition [Mateas 1999][Swartout et al. 2001][Cavazza *et al.* 2002][Cavazza *et al.* 2009], body gestures combined with speech [Cavazza *et al.* 2004][Dow *et al.* 2007][Cavazza *et al.* 2007][Lima *et al.* 2011a], hand-drawn sketches [Kuka *et al.* 2009][Lima *et al.* 2011b] and physiological inputs [Gilroy *et al.* 2012].

The use of natural language as a form of interaction in interactive storytelling systems was extensively explored by several authors. Façade [Matheas and Stern 2004] was the first large-scale use of natural language input in interactive storytelling. Façade accepts surface text utterances from the player and decides what reaction(s) the characters should have to the utterance. For example, if the player types "*Grace isn't telling the truth*", the system recognizes that this is a form of criticism and decides what reaction Grace and Trip should have after Grace being criticized. Its narrative is almost entirely dialogue-based and the high quality of the natural language system comes at the cost of significant authoring work [Mehta *et al.* 2007]. In a different approach, Mead *et al.* [2003] present a speech interface based on the notion of influence, where the user is able to interfere with the course of action by influencing the decisions made by the characters. This is achieved by recognizing corresponding speech acts and mapping them to the plans that implement characters' behaviors in the story. However, none of these works considers the possibilities of multi-user interactions.

Interactive TV is naturally a multi-user environment. In this context, Ursu *et al.* [2008b] propose new forms of interactivity with the TV content through the *ShapeShifting Media*, a system designed for the production and delivery of interactive screen-media narratives. Their applications include *My News & Sports My Way*, in which the content of a continuous presentation of news is combined in accordance with users' interest, and the romantic comedy *Accidental Lovers*, in which users can watch and influence a couple's relationship. In *Accidental Lovers*, viewers are able to influence the ongoing story by sending mobile text messages to the broadcast channel. Changes in the emotional state of the characters and their relationships depend on the existence of some specific keywords found in the viewer's messages. A similar approach is used in *Akvaario* [Pellinem 2000], where viewers can influence the mood of the protagonists through phone messages. There are few works on interactive cinema. *Last Call* [Jung von Matt 2010] is an interactive advert for the 13th Street TV Channel exhibited experimentally in movie theaters. In *Last Call*, the audience interacts with the actress talking to her via cell phone. However, none of these applications has a complete natural language interface that understands the user input; only specific keywords have some effect on the ongoing story. Furthermore, in *Last Call* the software simply works with a fixed tree of pre-recorded video segments.

The main difference between the interaction model presented in this work and the previous ones is the use of a natural language interface in a multi-user interactive environment for TV and cinema. Moreover, we adopt a fully implemented interactive storytelling system designed for interactive TV/cinema that automatically generates coherent and diversified interactive stories.

3. Interaction Model

Television and movies are classical media for telling stories. The possibility of interaction can expand the boundaries of those media towards new forms of storytelling. In this context, however, the need for comfortable and simple interaction methods becomes even more essential. Moreover, it is also important to consider not only the case in which users want to actively intervene in a story, but also the case in which they just want to watch TV without being called to interact by any means. Based on this idea, we established the following requirements for the proposed interaction interface:

- **Comfortable:** users must feel comfortable using the interaction interface.
- **Simple:** the interaction mechanism must be simple to allow anyone to interact easily.
- **Attractive:** the interaction interface must be attractive to motivate users to interact.

- **Transparent:** the interaction interface must be transparent to users that just want to watch TV/movies without interactions.
- **Multi-user:** the interaction interface must handle the multi-user interactions present in TV and movies.

In traditional interactive TV/cinema environments, interaction is usually done through on-screen GUI artifacts. However, there are two main problems with this approach: (1) it compromises the requirement for a transparent interface and may distract users that just want to watch TV/movies without interactions; and (2) it does not comply with the requirements of a complete multi-user interface, because on-screen GUI artifacts cannot cope with several users at the same room.

Our interaction model uses personal mobile devices (phones and tablets) to overcome the limitations of on-screen GUI artifacts and to provide a more attractive interaction mechanism. This is an effective approach because almost everyone has a mobile phone. Moreover, removing the interface artifacts from the TV/cinema screen makes the interaction more transparent and avoids interferences in the visual content being exhibited. As the mobile devices can be directly connected to the broadcast server through the internet, another great advantage of this model becomes visible: the interaction system is completely independent of the TV model; it is even compatible with traditional analog TVs.

The development of a user interface that is comfortable, simple, and attractive is another difficult challenge that emerges when designing an interaction interface for interactive storytelling. A simple GUI interface in the mobile devices may not be attractive and does not scale very well when there is a large amount of options to choose (which is very common in an interactive story). Our approach to create a more comfortable, simple, and attractive interaction mechanism is the use of a natural language user interface. In this system, users can write or speak what they want to happen in the story. With this interaction method it is easy to imagine the possibility of watching a movie while advices are continuously being sent to the characters.

3.1 Basic Architecture

Figure 3 illustrates the basic architecture of the natural language interface and how it communicates with the interactive storytelling system. The *Storytelling Server* generates and controls the execution of the stories, administrates multiple clients who share the same narrative, and informs the Interaction Server about valid suggestions for the next chapters. The *Interaction Server* is the connection between the interactive storytelling system and the multi-user interaction interface. It is responsible for receiving interaction messages, translating user's advices into valid story suggestions, and informing the Storytelling Server about the user's choices. The *Mobile Interface* is an application running on mobile devices that is responsible for capturing the user input (text or speech) and sending it through the internet to the Interaction Server. The *Client Drama Viewer* is responsible for the exhibition of the generated plots on the display device, which can be a TV or any other device than can receive and display the broadcast signal.

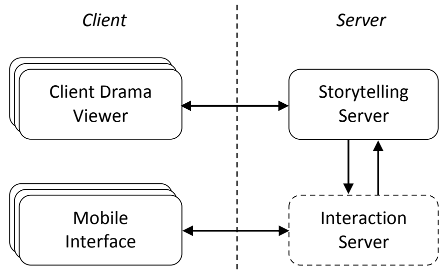

Figure 3: Basic architecture of the proposed multi-user natural language interaction system.

Users are able to interact by giving advices to the virtual characters at any time during dramatization. Figure 4 illustrates the dynamic behavior of the interaction system through an activity diagram. The Interaction Server is continuously collecting all the suggestions sent by the users (G facts) and combining them with the facts added (F^+) and removed (F^-) from the current state of the world by the planning system. When the end of the scene/chapter is reached, the facts that are more frequently mentioned by the users and that are not inconsistent with the ongoing story are then incorporated into the story plot. In Figure 4, thick black bars indicate parallel activities. We should notice that a chapter overture is exhibited (*e.g.* audio, text, and/or video) while the system runs the plot generator module (the most demanding processing time). Also we can see that interactions occur during the dramatization process (*i.e.* in parallel with the box "Run Drama" of Figure 4).

3.2 Mobile Interface

The Mobile Interface is the interface between users and the interaction system. It consists of a small application developed for Android mobile devices (such as smartphones and tablets), where users can interact with the ongoing stories by writing or speaking a suggestion/advice to the virtual characters using natural language. The user interface of this application is shown on Figure 5.

Users can interact with the story by typing the suggestions and advices in the text box shown on the mobile interface or by pressing the microphone icon and then speaking out the intended suggestion. The Android Speech Recognition API is used to recognize the user speech and to convert it into text. In this way, the system only needs to handle and understand written text. However, to make the mobile interface a very simple and light application, the process of extracting meaning from the user's suggestions is done in the Interaction Server. So, after reading the user input (text or speech), the mobile application sends the user suggestions through a TCP/IP connection to the Interaction Server, which is responsible for interpreting and extracting meaning from the suggestions.

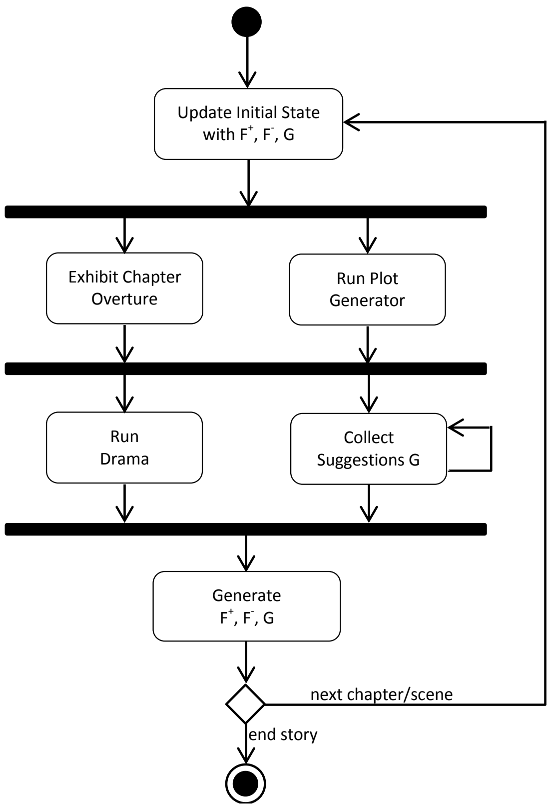

Figure 4: Activity diagram of the proposed interaction system.

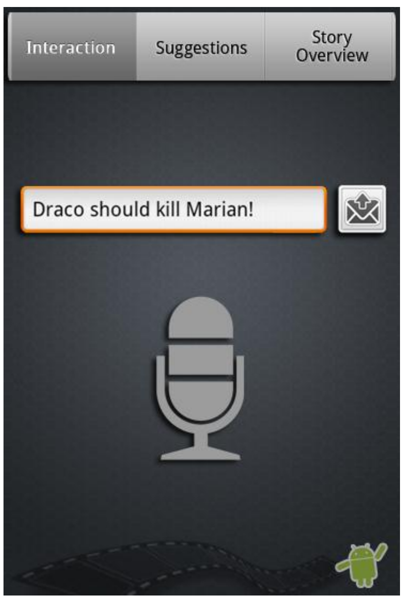

Figure 5: Mobile user interface.

4. Natural Language Processing

The process of extracting meaning from the users' suggestions involves natural language processing [Jurafsky and Martin 2000]. A traditional natural language processing task consists of two main phases: (1) syntax parsing, where the syntax tree and the grammatical relations between the parts of the sentence are extracted; and (2) semantic analysis, which is the extraction of the meaning of words or phrases.

In the proposed interaction system, we adopted the Stanford Parser to perform the syntax parsing of the sentences [Klein and Manning 2003]. The Stanford Parser [Stanford 2012] is a probabilistic parser that represents all sentence relationships as typed dependency relations instead of using phrase structure representations. However, it also produces phrase structure trees.

The Stanford Parser produces 55 different typed dependencies [Marneffe and Manning 2008]. These dependencies reflect the grammatical relationships between the words. Such grammatical relations provide an abstraction layer to the pure syntax tree and provide information about the syntactic role of all elements. Figure 6 shows a phrase structure tree generated by the Stanford Parser for the sentence "*Draco should kill Marian!*". The corresponding typed dependencies are listed in Figure 7. Typed dependencies facilitate the analysis of semantic relationships between words based on both their grammatical relationships and overall sentence syntactical structure.

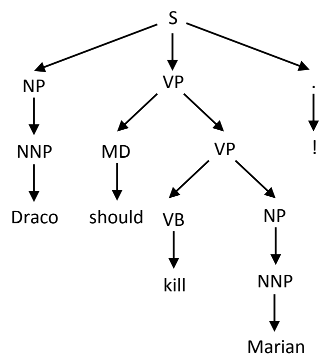

Figure 6: Phrase structure tree of "*Draco should kill Marian!*"

nsubj(kill-3, Draco-1)
aux(kill-3, should-2)
root(ROOT-0, kill-3)
dobj(kill-3, Marian-4)

Figure 7: Typed dependencies of "*Draco should kill Marian!*"

The typed dependencies are all binary relations, where a grammatical relation holds between a "governor" and a "dependent". In the above example, the relation *nsubj* (nominal subject) relates the noun "*Draco*" with the corresponding verb "*kill*", whereas the relation *dobj* (direct object) relates this verb with the object "*Marian*". In this way, the sentence elements are extracted and the sentence structure can be translated into simple first-order logic sentences. In the above example, the following sentence is extracted:

kill(*Draco, Marian*)

which means that "*Draco*" must perform the action "*kill*" and the victim is "*Marian*".

In the present work we generate simples logic sentences composed by a conjunction of predicates, e.g. from "*Brian and Hoel fight against Draco*" is generated the sentence "*fight*(*Brian, Draco*) and *fight*(*Hoel, Draco*)".

With this dependency chain, the system is able to extract "subject – direct object" relationships from sentences. However, for this pattern to be valid, four conditions must be met: (1) a nominal subject (*nsubj*) dependency must exist; (2) the dependent of the *nsubj* dependency must be a family member (in the phrase structure tree); (3) the governor of this dependency must be a verb, which means that a family member is the head noun of the subject of a clause which is predicated by the verb; and (4) a direct object (*dobj*) dependency must exist and the governor of this dependency must match the index of the governor of the *nsubj* dependency – then we assume that the dependency of the *dobj* relation is paired with the family member found initially.

In the example above, the extracted logical sentence already contains the semantic meaning needed by our interaction system to infer a valid suggestion to the story. However, there are some cases where the subjects are not directly referenced. For example, in the sentence "*Brian saves Marian and marries her.*", the pronoun "*her*" refers to "Marian". However, when we compute the typed dependencies for this sentence (Figure 8), we see in the relation "*dobj*(*marries-5, her-6*)" that the pronoun "*her*" was not resolved and, in some cases, it's not possible to solve it using only the phrase structure tree. The process of resolving what pronoun or a noun phrase refers to is called **anaphora resolution**. To solve this problem, we used another tool from the Stanford Natural Language Processing Group, the Stanford Deterministic Coreference Resolution System [Raghunathan et al. 2010], which is able to indicate precisely the correct reference of any unknown pronoun.

The parser also verifies the occurrence of negations. For example, in the sentence "*Draco should not kill Marian!*", the adverb "*not*" completely changes the meaning of the sentence. To identify negations, the parser analyses the occurrence of negation modifiers ("*neg*") in the typed dependency list. Figure 9 illustrates the typed dependency for the example above and the occurrence of the negation modifier.

nsubj(saves-2, Brian-1)
nsubj(marries-5, Brian-1)
root(ROOT-0, saves-2)
dobj(saves-2, Marian-3)
conj_and(saves-2, marry-5)
dobj(marries-5, her-6)

Figure 8: Example of anaphora problem in the sentence "*Brian saves Marian and marries her.*"

nsubj(kill-4, Draco-1)
aux(kill-4, should-2)
neg(kill-4, not-3)
root(ROOT-0, kill-4)
dobj(kill-4, Marian-5)

Figure 9: Example of negation in the sentence "*Draco should not kill Marian!*"

After translating the "*subject – direct object*" relations into first-order logic sentences, the parser also needs to validate the sentences. For example, the predicate "*fight*(CH_1, CH_2)" requires a nominal subject CH_1 that is a valid character and a direct object CH_2 that also is a valid character in the story context. Moreover, the verb "*fight*" also must be a valid action. To perform this validation, the parser has access to a list of valid actions, characters and places. In this way, the parser is able to identify the elements that the words represent. However, almost all words have synonyms and to deal with this, the parser also incorporates a dictionary of synonyms associated with each action, character and place. So, it is able to parse sentences such as "*The hero should annihilate the villain!",* where the verb "*annihilate*" is a synonym of the action "*kill*", and the objects "*hero*" and "*villain*" are, respectively, the roles of the characters "*Brian*" and "*Draco*".

Ideally, the parser expects sentences that contain at least one verb, one nominal subject and a direct object. However, this not always happens; in some cases the subject, the direct object, or both are omitted. For example, the sentence "*Kill the princess!*" does not express directly who should perform the action "*kill*", but indicates the direct object "*princess*" (Figure 10). In this case, the parser is still able to generate a partial logic sentence to represent it:

kill(*, *Marian*)

which means that someone "*" must perform the action "*kill*" and the victim is "*Marian*" (identified by its role in the story ("*princess*")). The operator "*" can be replaced by any valid character to complete the logical sentence.

```
root(ROOT-0, Kill-1)
det(princess-3, the-2)
dobj(Kill-1, princess-3)
```

Figure 10: Example of omitted subject in the sentence "*Kill the princess!*"

The entire process of extracting valid first-order logic sentences from text phrases is illustrated in Figure 11. In the Syntax Parsing step, the Stanford Parser receives a text phrase S_x as input and generates a Dependency Tree and the Typed Dependencies for the sentence. Using this information, in the Semantic Analysis phase, the parser performs the Anaphora Resolution process to resolve the pronouns of the sentence and find valid synonyms using the Synonym Dictionary. Finally, the parser checks the integrity of the sentences using some Logic Rules and returns a list of valid first-order logic sentences (P_x^n).

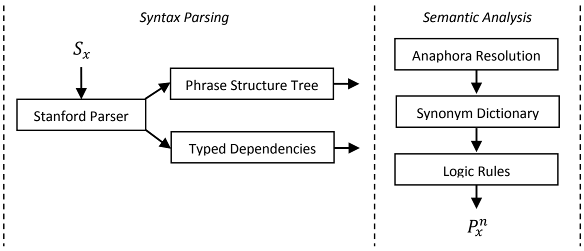

Figure 11: The process of extracting valid first-order logic sentences. S_x is the input text phrase and P_x^n is the output list of predicates

5. Evaluation

To evaluate the multi-user natural language interface, we performed a technical test to check the performance and accuracy of the proposed method of interaction. The evaluation was based on two experiments: (1) the *recognition rate test*, to check the accuracy of the predicted suggestions; and (2) the *performance test*, to check the time needed to process the input comments and recognize the suggestions as first-order logic sentences. For both tests, we used a set of 107 text suggestions collected from users that were testing the interaction system. After a manual analysis of the suggestions, we found 81 suggestions that were manually classified as valid suggestions.

For the recognition rate test, we used our method to extract valid story suggestions from the text suggestions and then compared the results with the results obtained by the manual classification. As a result we got a recognition rate of 90.6%, with only 10 valid comments being incorrectly classified as invalid suggestions. The main reason for the incorrect classifications was the occurrence of spelling mistakes in the comments.

To evaluate the performance of our method, we again utilized the collection of 107 comments collected from users, and calculated the average time necessary to perform the recognition of the suggestions as first-order logic sentences. The computer used to run the experiments was an Intel Xeon E5620, 2.40 GHZ CPU, 24 GB of RAM using a single core to process the algorithms. As a result we got the average time of 2.7 milliseconds to process an input comment and recognize the suggestion as first-order logic sentences (standard deviation of 1.3 milliseconds).

In our experiments, the multi-user natural language interface produced good results. However, natural language processing is not a trivial task. It is possible that our parser will not correctly recognize every possible valid sentence, but we believe that it will be able to recognize the sentences in the most part of the cases without the audience being aware of mistakes. The time necessary to process the user suggestions is small, but it grows according to the number of suggestions to be processed. With a large number of users interacting at same time, the parallelization of this process may be necessary to guarantee the real-time execution of the narrative.

6. Conclusions

In this paper, we presented a multi-user natural language interface for interactive storytelling applications designed for TV and cinema. The proposed interaction method allows users to watch an interactive movie and give suggestions/advices to virtual characters through a mobile device – and then observe the advices actually affecting the decision made by the characters. As far as we are aware, this is the first time this form of interaction is explored in a fully implemented interactive narrative designed for TV and cinema.

The prototype was built over the Logtell system; however its architecture is generic enough to be adopted by any interactive storytelling system organized by chapters or scenes. Moreover, the interaction mechanism can be adapted to other systems that require a multi-user natural language interface.

As future works we intend to improve our interaction system by including a module to automatically correct spelling mistakes in the user suggestions, which will probably increase the recognition rates. Also we intend to conduct a user study to evaluate the system usability from a Human-Computer Interaction (HCI) perspective and compare the proposed interaction mechanism with other forms of interaction in interactive narratives.

Acknowledgements

This work was partially supported by CAPES (Coordination for the Improvement of Higher Education Personnel, linked to the Ministry of Education) under grant RH-TVD 01/2007 N°. 133/2008, and by FINEP (Brazilian Innovation Agency) and CNPq (National Council for Scientific and Technological Development) belonging to the Ministry of Science and Technology.

References

CAVAZZA, M., CHARLES, F. AND MEAD, S., 2002. Character-based interactive storytelling. *IEEE Intelligent Systems, Special issue on AI in Interactive Entertainment*, vol. 17, (4), 17-24.

CAVAZZA, M., CHARLES, F., MEAD, S. J., MARTIN, O., MARICHAL, X. AND NANDI A., 2004. Multimodal acting in mixed reality interactive storytelling. *IEEE Multimedia*, vol. 11, (3), 30-39.

CAVAZZA, M., LUGRIN, J-L., PIZZI, D. AND CHARLES, F., 2007. Madame bovary on the holodeck: immersive interactive storytelling. *In: Proceedings of the 15th International Conference on Multimedia (MULTIMEDIA'07)*, 651-660.

CAVAZZA, M., PIZZI, D., CHARLES, F., VOGT, T. AND ANDRÉ, E., 2009. Emotional Input for Character-based Interactive Storytelling. *In: Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, Budapest, Hungary, 313-320.

CHESHIRE, T. AND BURTON, C., 2010. *Transmedia: Entertainment reimagined* [online], Wired UK, Available from: http://www.wired.co.uk/magazine/archive/2010/08/features/what-is-transmedia [Accessed 20 April 2012].

CIARLINI, A.E.M., POZZER, C.T., FURTADO, A.L. AND FEIJO, B., 2005. A logic-based tool for interactive generation and dramatization of stories. *In: Proceedings of the International Conference on Advances in Computer Entertainment Technology*, Valencia, 133-140.

DOW, S., MEHTA, M., HARMON, E., MACINTYRE, B., AND MATEAS, M., 2007. Presence and Engagement in an Interactive Drama. *In: Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI'07)*, San Jose, 1475-1484.

GILROY, S., PORTEOUS, J., CHARLES, F., AND CAVAZZA, M., 2012. Exploring Passive User Interaction for Adaptive Narratives. *In: Proceedings of the International Conference on Intelligent User Inter-faces (IUI)*, ACM, 119-128.

GRASBON, D., BRAUN, N., 2001. A morphological approach to interactive storytelling. *In: Proceedings of Cast01, Living in Mixed Realities*, Sankt Augustin, Germany, 337-340.

JUNG VON MATT/SPREE, 2010. Last Call, Berlin, Available from www.youtube.com/watch?v=qe9CiKnrS1w [Accessed 01 August 2012].

JURAFSKY, D. AND MARTIN, J. H., 2000. Speech and language processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Prentice-Hall.

KLEIN, D. AND MANNING, D.C., 2003. Accurate Unlexicalized Parsing. *In: Proceedings of the 41st Meeting of the Association for Computational Linguistics*, 423-430.

KUKA, D., ELIAS, O., MARTINS, R., LINDINGER, C., PRAMBÖCK, A., JALSOVEC, A. MARESCH, P., HÖRT-NER, H. AND BRANDL, P., 2009. DEEP SPACE: High Resolution VR Platform for Multi-user Interactive Narratives. *In: Proceedings of the 2nd Joint International Conference on Interactive Digital Storytelling*, 185-196.

LIMA, E.S., FEIJÓ, B., BARBOSA, S., SILVA, F.G., FURTADO, A.L., POZZER, C.T. AND CIARLINI, A.E.M., 2011A. Multimodal, Multi-User and Adaptive Interaction for Interactive Storytelling Applications. *In: Proceedings of the 10th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, Salvador, Brazil.

LIMA, E.S., FEIJÓ, B., BARBOSA, S.D.J., FURTADO, A. L., CIARLINI, A. E. M. AND POZZER, C.T., 2011B. Draw Your Own Story: Paper and Pencil Interactive Storytelling. *In: Proceedings of the 10th International Conference on Entertainment Computing (ICEC)*, Lecture Notes in Computer Science, Vol. 6972, 1-12.

LIMA, E.S., FEIJÓ, B., FURTADO, A.L., POZZER, C., CIARLINI, A., 2012. Automatic Video Editing For Video-Based Interactive Storytelling. *In: Proceedings of the 2012 IEEE International Conference on Multimedia and Expo (ICME)*, Melbourne, Australia, pp. 806-811.

LOGTELL, 2012. The Logtell Project Website, Available from: http://www.icad.puc-rio.br/~logtell [accessed 19 April 2012].

LOYALL, A.B. AND BATES, J., 1991. *Hap: A Reactive, Adaptive Architecture for Agents.* Technical Report CMU-CS-91-147, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

MARNEFFE, M. AND MANNING, C.D., 2008. The Stanford typed dependencies representation. *In: Proceedings of the Workshop on Cross-Framework and Cross-Domain Parser Evaluation*, Manchester, 1-8.

MATEAS, M. AND STERN, A., 2003. Façade: An experiment in building a fully-realized interactive drama. In: Game Developers Conference, Game Design track [also available in www.interactivestory.net/papers/MateasSternGDC03.pdf].

MATEAS, M. AND STERN, A., 2004. Natural Language Understanding in Façade: Surface-text Processing. *In: Proceedings of Technologies for Interactive Digital Storytelling and Entertainment (TIDSE)*, 3-13.

MATEAS, M., 1999. An Oz-Centric Review of Interactive Drama and Believable Agents. *In: AI Today: Recent Trends and Developments*, Wooldridge, M., Veloso, M. (eds.), Lecture Notes in AI, vol. 1600, Springer Verlag, New York, 297-328.

MEAD, S.J., CAVAZZA, M. AND CHARLES, F., 2003. Influential Words: Natural Language in Interactive Storytelling. *In: Proceedings of the 10th International Conference on Human-Computer Interaction*, Crete, Greece, 741-745.

MEEHAN, J.R., 1977. TALE-SPIN, An Interactive Program that Writes Stories. *In: Proceedings of the 5th International Joint Conference on Artificial Intelligence*, 91-98.

MEHTA, M., DOW, S., MATEAS, M. AND MACINTYRE, B., 2007. Evaluating a Conversation-Centered Interactive drama. *In: Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007),* Honolulu, Hawaii.

PELLINEN, T., 2000. Akvaario (Aquarium), Media Lab, Helsinki University of Art and Design Finland. Broadcast by The Finnish Broadcasting Company.

RAGHUNATHAN, K., LEE, H., RANGARAJAN, S., CHAMBERS, N., SURDEANU, M., JURAFSKY, D. AND MANNING, C., 2010. A Multi-Pass Sieve for Coreference Resolution. *In: Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing,* Boston, 492-501.

STANFORD, 2012. The Stanford NLP (Natural Language Processing) Group. Available from: http://nlp.stanford.edu/software/lex-parser.shtml [accessed 20 April 2012].

SWARTOUT, W. ET AL., 2001. Toward the Holodeck: Integrating Graphics, Sound, Character and Story. *In: Proceedings of the 5th international conference on Autonomous Agents*, ACM Press, New York, 409-416.

URSU, M.F., KEGEL, I.C., WILLIAMS, D., THOMAS, M., MAYER, H., ZSOMBORI, V., UOMOLA, M.L., LARSSON, H. AND WYVER, J., 2008b. ShapeShifting TV: Interactive screen media narratives. *Multimedia Systems*, 14, (2), 115-132.

URSU, M.F., THOMAS, M., KEGEL, I.C., WILLIAMS, D., TUOMOLA, M., LINDSTEDT, I., WRIGHT, T., LEURDIJK, A., ZSOMBORI, V., SUSSNER, J., MYRESTAM, U., AND HALL, N., 2008a. Interactive TV Narratives: Opportunities, Progress, and Challenges. *ACM Transactions on Multimedia Computing, Communications and Applications*, 4 (4), Article 25.

WILLIAMS, D., URSU, M.F., MEENOWA, J., CESAR, P., KEGEL, I. AND BERGSTRÖM, K., 2011. Video mediated social interaction between groups: System requirements and technology challenges. *Telematics and Informatics*, 28, 251-270.