Towards Virtual Communities on the Web: Actors and Audience

Anton Nijholt (Parlevink Research Group)[1]
University of Twente, PO Box 217
7500 AE Enschede, the Netherlands
anijholt@cs.utwente.nl

Abstract: *We report about ongoing research in a virtual reality environment where visitors can interact with agents that help them to obtain information, to perform certain transactions and to collaborate with them in order to get some tasks done. Our environment models a theatre in our hometown. We discuss attempts to let this environment evolve into a theatre community where we do not only have goal-directed visitors, but also visitors that that are not sure whether they want to buy or just want information or visitors who just want to look around. It is shown that we need a multi-user and multi-agent environment to realize our goals. Since our environment models a theatre it is also interesting to investigate the roles of performers and audience in this environment. For that reason we discuss capabilities and personalities of agents. Some notes on the historical development of networked communities are included.*

1 Virtual Community Development

We discuss a virtual reality theatre environment in which we have embedded agents that can help the user through natural language dialogue. The environment has been built using VRML (Virtual Reality Modeling Language) and is accessible from all over the world on WWW. In [5] we discussed a natural language dialogue system that offered information about performances in some of our local theatres and that allowed visitors to make reservations for these performances. This dialogue system has been assigned to a visualized embodied agent to which users can ask questions. In addition, other agents have been defined, allowing other (primitive) dialogues with users or visitors of the system. We discuss how our ideas about this system changed in time by adding more facilities to it and by paying more attention to potential users. Rather than a goal-directed information and transaction system, the environment is now evolving into a virtual community where differences between visitors and artificial agents become blurred.

Before going into details of our own environment we think it is useful to zoom in on the historical development of networked communities.

1.1 Text-based Virtual Worlds

The first networked virtual worlds were text-based. They became known as MUDs (Multi-User Domains). and allowed synchronous communication between users and access to a shared database with text descriptions of users and objects. These worlds were designed as multi-user text-based games where users not only communicate but also collaborate or fight one another. In later years 'social' MUDs appeared with the emphasis on social interaction and also, stimulated by the advent of MOOs (Multi-user Object Oriented Environments), for educational purposes. Internet Relay Chat (IRC) technology allowed further development of social and recreational text-based computer mediated communication.

In these text-based environments the personality of a user, as it appears to others, is based on the contents and the style of the text utterances the user produces in the dialogues, his or her turn taking behavior (aggressive, cooperative, shy, ...) and more generally the moods (as they show) and attitudes towards different participants and to the community that can develop in such environments. But also a person’s name has connotations related to personality. Obviously, all kinds of otherwise probably irrelevant details (typing speed, quality of the connection, . .) can influence other's judgments. Moreover, there is the possibility to 'cheat' by pretending to be someone else, with a different personality, gender and physical appearance or even a different mood than in real life at that moment. Some environments have developed into role-playing chat communities where users want to live a different life in a different body and with a different personality.

1.2 2D-Graphical Virtual Worlds

Graphical multi-user environments were introduced in the 1980s. In a typical setting we have a background image showing the entrances to several locations or rooms in the environment or we are already in one of these 2D locations and we can choose one of the other visitors (or all of them) to talk to. Typically, visitors can present themselves by choosing an avatar (a 2D object) and its predefined animations (express joy, sadness, . .).

[1] Research reported in this paper has been made possible by the “VR Valley Twente” foundation and by the U-Wish project of the Dutch Telematics Institute.

Animations are simple: a waving gesture, a jump of joy, . . .). Interactions with users take place using these animations to express certain emotions. Most interactions are text-based, by using chat windows and text balloons that appear above the head of avatars that take part in the discussion. A well-known 2D graphical chat world is The Palace (Damer [2]). While a visitor of these worlds can select and move an avatar which may personify him or fit the role he plays, it seems that nevertheless in textual worlds visitors have much longer interactions, get more personal and build up deeper and longer relationships. On the other hand, rather than having discussions and chats in these worlds visitors can much easier organize events in locations where visualization supports the discussions that take place.

1.3 3D-Graphical and Virtual Reality Worlds

With the advent of VRML 3D virtual worlds could be designed for Worldwide Web. Special browsers are needed to visit and explore these worlds (Cosmo Browser, Sony's Community Player). Not all VRML browsers allow multiple users to share an environment. Rather than for chatting, the worlds were meant to be explored, to explain or to allow the simulation of a particular activity in which the visitor had to be involved. Virtual reality applications were already there and rather than consider distributed virtual reality as a technology to design communities it was explored for all kinds of applications. Especially in educational applications the virtual environment is considered as a constructionist environment. That is, rather than learning by being told, the student learns by creating. Nevertheless, virtual worlds intended to meet other people entered the arena. Some well-known examples are Active Worlds, Worlds Chat, OnLive!Traveler and Blaxxun's Colony City. In these worlds multiple users can share the scenes. The users are represented as avatars and they can chat with each other using chat windows.

These virtual worlds are not necessarily completely preconstructed. Users can extend and change the world by building homes and environments on unoccupied area. Different features for seeing and communicating can be offered to users. E.g., it is possible to 'whisper' to a user in close geographical proximity, to 'mute' another user, to send telegrams to other users and to set privacy levels. Conversations between multiple users (with text-based dialogues) will have different threads, which are not always easy to follow.

1.4 Using 3D Avatars

In the more advanced worlds users can have more sophisticated visual representations that can be animated and that can interact with each other without being restricted to a very limited set of predefined gestures. As an example, OnLive! Traveler allows 'talking heads' as avatars. A user can choose a 3D head and customize it in a limited way. The heads have some facial gestures, but more importantly, the OnLive! VRML browser allows real-time voice communication between multiple users, where this communication is accompanied by lip-synchronization of the heads. Users whose heads are near the head of a speaker will all hear the speaker's voice. Hence, a group conversation is possible, but there are also options to have private conversations.

In Active Worlds, to mention an other example, avatars have a name and can express gender, culture, age and they can display gestures, emotions and actions. The personality of the user can be reflected by these choices, as it can be reflected by the homes and environments the user builds. The animations (sometimes including facial gestures) are pre-programmed. They can be compared with the emoticons in text-based communication. That is, rather than including a ☺ in the text, clicking a button of the browser makes an avatar jump up and down to express that the user is happy. Users can select 3D avatars from a library, customize them, and create them using an avatar wizard. Users can design and import their own 3D avatars, but this requires a basic understanding of 3D graphics (in Active Worlds), VRML (in DeepMatrix and Blaxxun) and animation. An avatar can be made to resemble the human user by photographic means.

Talking about privileges, in some communities certain privileges accompany (paid) registration. These privileges allow to select or construct a better quality of avatars, to build or claim property, to take part in the economic and governance system and have more freedom and privacy in communication with other community members.

For its user the avatar is in fact the camera by which the scene is viewed. This is the first person perspective. In the orthographic view (the third person perspective) users can see their visual representation in the virtual world. Blaxxun, for example, allows users to toggle between these perspectives.

1.5 From 3D Avatars to Embodied Agents

In the communities described so far inhabitants had the possibility to represent themselves by their writings (in text-based virtual worlds) or by choosing a 2D or 3D avatar in combination with chatting with other community members. Simple animations help to express themselves and users can move their representations from one context to an other context. The worlds that we considered may have collision and gravity features which may become visible in the movements of avatars or by audible feedback. Because of the real-time voice communication, in OnLive! Traveler there is no need to shift the focus of attention between a VR window and a chat window. In addition we can recognize the users voices and look at lip-sync facial gestures. These features have been reported to be responsible for a high level of the sense of being embodied and of being present in the OnLive! environment.

There is an enormous gap when we compare the capabilities of the avatars and talking heads with those of the humans they represent. One way to close this gap is to give the human user the ability to control the avatar in a much more detailed way. This requires sophisticated avatars and animations explicitly controlled by the user or captured from verbal and non-verbal input or from body movements of the user. In addition to the avatars that represent humans we can also add avatars to the environment to increase the sense of reality. They should at least be animated, but preferably there should be possibilities to give them personality and capabilities to act on their own or on behalf of their users. They need appropriate internal modeling. Agents are needed.

Generally, an agent is situated in an environment in which it is capable of autonomous action. Agents are intelligent when they can respond on changes in their environment, when they can take the initiative and when they are able to interact with other agents and humans to solve problems. After having discussed our environment in section 2 we will discuss properties of agents in this environment in sections 3 and 4.

2 Building a Theatre Environment

Our starting point in designing a virtual theatre environment was the situation described in the previous subsection. The theatre was built according to design drawings of the architects of the building. Visitors can explore this environment, walk from one location to another, ask questions to available agents, click on objects, etc. Karin (see Fig. 1), the receptionist of the theatre, has a 3-D face that allows simple facial expressions and lip movements that synchronize with a text-to-speech system that mouths the system's utterances to the user. For the lips we have five visemes, responsible for some fifty different Dutch phonemes. For the generation of system utterances templates are used that can be annotated with intonation and facial animation information. This has not been done yet. Design considerations that allow an embodied agent like Karin to display combinations of verbal and non-verbal behavior can be found in [9]. Because of web limitations, there is no sophisticated synchronization between the (contents of the) utterances produced by the dialogue manager and corresponding lip movements and facial expressions of the Karin agent.

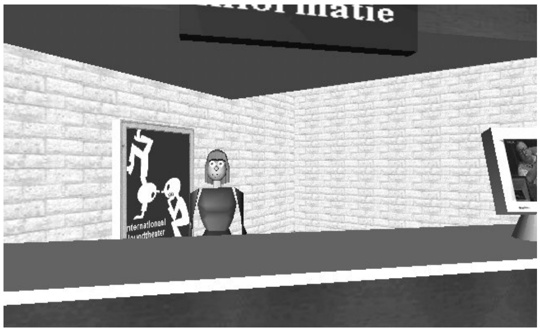

Fig 1 Karin behind the Information Desk

Other agents in this environment have been introduced. One example is a navigation agent, which knows about the building and can be addressed using speech and keyboard input of natural language. No real dialogues are involved. The visitor can ask about existing locations in the theatre and when recognized a route is computed and the visitor's viewpoint is guided along this route to the destination. This navigation agent has not been visualized as an avatar. Its viewpoint in the theatre is the current viewpoint from the position (coordinates) of the visitor in the world. A Java based agent framework has been introduced to provide the protocol for communication between agents. It allows the introduction of other agents. For example, why not allow the visitor to talk to the theatre seat map or to a poster displaying an interesting performance? This has indeed been made possible. Unlike a predecessor, the version of the virtual theatre with a speech recognizing navigation agent has not been made accessible to the general audience by putting it on the Web. Although speech recognition is done at the server (avoiding problems of download time, ownership, etc.) there are nevertheless too many problems with recognition quality and synchronization with the events in the system. However, further work on the navigation agent is in progress. Part of this work is on user preferences for navigation in virtual worlds, part is on modeling navigation knowledge and navigation dialogues, part is on adding instruction models to agents and part is on visualization.

3 Multi-agents and Multi-users

In our environment we can have different human-like agents. Some of them are represented as communicative humanoids, more or less naturally visualized avatars standing or moving around in the virtual world and allowing interaction with visitors of the environment. In a browser that allows the visualization of multiple users, other visitors become visible as avatars. We want any visitor to be able to communicate with agents and other visitors, whether visualized or not, in his or her view. That means we can have conversations between agents, between visitors, and between visitors and agents.

In the previous sections we talked about agents acting in our own virtual theatre. Karin was introduced as a 'visualization' of our existing dialogue system. She has extensive knowledge of performances that play in the theatre. She can move her lips and have some simple head movements in function of the dialogue. Once we had Karin it became clear that we needed an agent framework and in it we introduced a navigation agent with some geographical knowledge and speech recognition capabilities. In fact, we have a multitude of potential agents. For example, we have a piano player on stage with some simple predefined animations accompanying the music. At the Università degli Studi di Milano research has been done on baroque dance animation with virtual dancers [1]. Using a baroque dance editor dances performed by virtual dancers can be choreographed and generated. Since the generated dances and animations are described in VRML it has become possible to have some guest performances of the Scala of Milan dancers in our theatre. In Fig 2 we see a visitor's avatar (Jacob) of the multi-user environment that has been so impertinent to climb the stage in order to get a closer look at the performing dancer. Its animations allow it to walk around following the coordinates of the moving viewpoint position of its owner.

To maintain a virtual environment where we have a multitude of domain and user-defined agents we need some uniformity from which we can diverge in several directions and combinations of directions: agent intelligence, agent interaction capabilities, agent visualization and agent animation. This will be discussed next.

We can look at some VRML related standards that have been proposed or are under development. Although we expect to make a transition from VRML to Java3D in the near future, similar standards will emerge then. One of the present standards is the Humanoid Animation (H-Anim) standard. It defines a structure and interface for (animated) avatars in VRML. The Living Worlds Standard aim is to define a conceptual framework and interfaces to support the creation of multi-user and multi-developer applications in VRML. They allow that applications can be assembled from libraries of components developed by multiple suppliers. Living Worlds deals with data distribution and scene synchronization. At a lower level there are standards on network and application protocols. At a higher level are issues that have to be dealt with when introducing interacting agent frameworks in virtual environments.

The visitor's avatar shown in Fig. 2 has been built following the H-Anim standard. Presently we use the DeepMatrix [14] multi-user environment system which is compliant with the Living Worlds specification.

In order to develop our environment in a direction where we can have interactions between the agents that

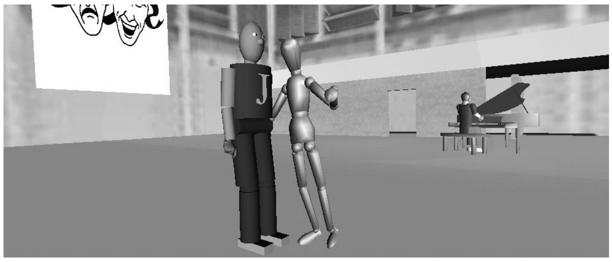

Fig. 2 Visitor, Dancer and Piano Player on the Virtual Stage

perform different tasks, agents that interact in multimodal ways with users and agents that interact with users avatars we need to be able to provide agents with different capabilities and personalities. The following lines of research have to be taken in order to be able to maintain a complex shared environment with interacting agents. First, we need to redesign and extend our agent framework such that individual agents both can represent (human) visitors (e.g., their movements, posture, nonverbal behavior) and can stand for domain agents. Second, visualization and animation of agents should comply to standards. Agents should be able to know (and reason) about physical appearance and their animation capabilities. Third, the internal modeling of agents should fit in agent frameworks that pay attention to issues of autonomy, reactivity, pro-activity, social ability and learning, in order to profit from developments in agent technology and to allow interoperability. It may be useful to be able to put every agent somewhere in the range going from completely controlled (by a user) to completely autonomous.

Next we discuss agents when they take their roles as actors and audience in our virtual theatre environment.

4 The Stage: Actors and Audience

In this section we want to concentrate on the interaction issues that play a role when designing an interactive virtual theatre community. As mentioned by Laurel [6], the computer screen can be looked upon as a stage and it has been argued that this theater metaphor can help in understanding and improving human-computer interaction. Allowing our environment to be used as a theatre community, including facilities for theatre professionals, but most of all for people interested in theatre and willing to experiment with others by staging performances themselves, will give us the opportunity to design and study all kinds of advanced HCI issues.

4.1 Stakeholders and Theatre

In the previous sections, we had traditional views on:

- the public, that is, men, women, children who want attend a certain performance or who want to know about performances in general, in a certain city or region, and at a certain date or in a certain period; the public has expectations about the information that is provided, it knows, for example, that different newspapers have different opinions about performances, hence, it is necessary to be careful in pushing the visitor to attend a certain performance

- the theatres, that is, the organizations that sell tickets, want positive reviews for their performances, want to give correct and relevant information to the public, and offer contracts to managers, artists and theatre companies in such a way that they are not loss-making

Now that we have a virtual theatre where people can look around and get information on performances, can we make the picture complete and apply this virtual reality environment to other theater-related purposes?

Apart from offering facilities and anticipating behavior of visitors interested in (buying tickets for) performances in our environment we can look more closely at possibilities that can be offered to:

- the professionals (stage directors, choreographers, stage crew, sound and light people, etc.)
- the performers, hence, the actors, the musicians, the dancers, the artists, authors and poets who present their work and prefer more or other interaction with each other or/and the audience
- the public in its role of audience attending a performance; not necessarily passive, just enjoying, but also as web-audience that can (real-time) influence the course of actions during a performance, take part in a performance by taking the role of an actor or having tools available to stage their own performances with friends and relatives while they are not necessarily present in the same location

Here we will not elaborate the possibility to use our environment for scenographic simulations. There are projects aiming at providing tools and environments to help in pre-producing performances. In these projects users can build a scenography of a performance, they can move through virtual models of stage sets in real time, they can experiment with lights or camera effects, change points of view, and preview a performance using animated human figures. However, better facilities for performance design can be offered on specialized hard- and software. Therefore, we concentrate on the possibility to look at our environment as a stage on which we can have on-line performances.

4.2 Actors and Theatre

When we have (distributed) performances and performers, we can distinguish

- 'Canned' avatar and 'streaming' video-audio performances; no control by or interaction with the on-line audience, although it might be possible that a visitor can 'participate' by being allowed to take the viewpoint of an avatar during a performance.
- Avatars representing players or members of the audience. The movements of the avatars can be obtained from the movements and manipulations of the players (using motion tracking sensors, haptic devices, interactive gloves, speech commands or (3D) mouse and keyboard input). There is not necessarily a one-to-one mapping of movements; they can be amplified or converted to other actions of the avatars; in fact, the avatars can be more or less obedient and may have some (semi-) autonomous (built-in) behavior even when they represent and have been introduced by a human actor.
- 'Autonomous' theatre avatars that have a particular role (on stage or in the audience) but that nevertheless have intelligence that allows them to react on other theatre avatars and, of course, visitors that are represented by avatars, as described above.

In the past, different types of performances in virtual reality on WWW have been held, including Shakespeare's *A Midsummer Night's Dream* and *Romeo and Juliet* with synthesized and captured movements.

4.3 Audience and Theatre

Let us now turn to the role of the audience. In the traditional theatre, performers and audience are physically together. There is a focus of attention of the audience in things happening on stage and performers are aware of the audience's attention. Rather than to have one special physical space where performers and audience gather, now performers can be geographically dispersed and so can the audience. Moreover, there is no need to maintain the distinction between audience and performers. The environment should allow an (web) audience that can (real-time) influence the running of things during a performance or can even take part in a performance by taking the role of an actor. This requires special attention for the presence issue, both for actors and audience (Reeve [13]).

A distinction between direct interaction (a user joins an ongoing performance as one of the actors or players in it) and ritualistic interaction has been made in Sgouros [16]. The latter form of interaction influences the performance of an event (e.g., a football match) but it also unifies participants and mobilizes emotions and sentiment. The web system that is given has a server which relays messages by the participants (players and audience). The audience can send messages to express emotions, to show approval of player actions, to warn a player, etc., using pre-defined and free-form messages with corresponding audio and image effects. The server can detect audience preferences and express collective reactions. Obviously, this may also allow the server to influence the performance of the event.

The Sgouros system shows some principles - and the possibilities to implement them - of different forms of audience interactions. Do we want to represent the audience? When we visualize the audience its reactions to a performance can be made visible too. Moreover, we can try to simulate crowd motion in order to have the audience enter the performance hall, take their seats and cheer and boo during a performance. This may improve the sense of participation for an individual viewer sitting lonely at home. When this participant represents itself in the theatre as a (smart) avatar) this avatar has to interact with the crowd. In this crowd also some individualities can be generated. Anyway, the visitor's avatar needs to match its motions with that of others, keep individual space in order to avoid collisions, find its seat among others doing the same, etc. (cf. Musse & Thalmann [8]). Fig. 3 shows an example of the work of Musse.

Crowd modeling is also interesting when we have a group of actors performing on stage, all of them, with slight variations, doing similar movements as in ballet. An individual avatar, more or less controlled by a human actor, may interact with such a group. On the other hand, even the crowd behavior may be given under control of a human actor. Just as in the case of individual actors and their body movements, group movements in virtual reality can express feelings and emotions (Price et al [11]).

As mentioned in Reitmayr et al. [14], having a large crowd as audience introduces all the real world logistics of event presentation, including seat assignment and sight lines. An alternative would be to enter a performance hall without having to share it with others, being able to take the best seat, to move on-stage (as in Fig. 2), etc. Again, audience feedback could be registered as metadata and returned as simulated applause or shouts of disapproval to each audience member.

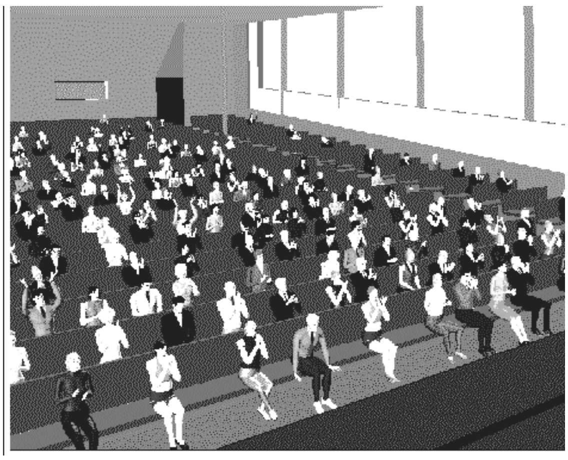

Fig. 3 Audience modeling at MIRALab, Geneva

5 Interactions and Their Control

It is not enough – on stage, during joint work or joint recreation – to be able to speak to set up and maintain a conversation. In order to model naturalistic group interaction behavior body movements, head pose, gaze direction, facial expressions, gestures and prosody are issues that should be looked at. How do we model or determine the focus of attention of an agent when it has the possibility to switch its attention from one event to another, from one agent to another? For the development of our environment we would like to see automatic or controlled generation of naturalistic group interaction behavior (with backchannels and for two or more agents) and at least be able to interpret what's going on during group interaction behavior in order to give individual agents this capability or to support a human user in its interactions. In cognitive research, psychology and linguistics attention has been given to these issues. When in combination with developments in computer science computer-supported collaborative work, video- and audio conferencing, virtual class and meeting rooms were studied, systems to support natural user interaction with the aim to make it more effective were developed.

Research that helps us to go in this direction is described in Rousseau et al. [15] (mixed initiative in agent interaction), Das et al. [3] (amplifying human perception and cognition) and Monzani et al. [7] (simulating interaction behavior using virtual sound propagation). One of the issues we are experimenting with is gaze modeling. It is assumed that in future implementations of our system gaze will not only play a role as input and output modality (e.g., to regulate the conversation between a user and a system's agent), but also in conversations between a user and more than one agent in the system. To gain experience with the gaze modality, we are implementing findings on gaze in a separate prototype environment with an eye tracking system that establishes where a user looks at [17]. Results of the experiments can also be used in the interaction between agents, e.g., on the stage.

6 Conclusions

Our research on agents and avatars in virtual worlds and our ideas how to develop such worlds as communities have been seriously influenced by Neil Stephenson's novel *Snow Crash*. In this book there are two equally existing and equally important worlds: *Reality*, the real world and a virtual world called *Metaverse*. People can buy or rent real estate and build houses or hire people to build houses in Metaverse. Some parts of this world are expensive to live, other parts can be afforded by the poor. People who don't have a house in Metaverse can make use of public terminals in order to enter the main street of Metaverse. Avatars representing people coming from public terminals can be recognized since they are 'trashy' and black-and-white. Avatars can walk around, but there is also public transport and the privileged drive cars or motor-cycles, can take passengers and can have races. Avatars sometimes take an elevator, they can fight (e.g., in free-combat zones), smoke cigarettes (where the rendering of smoke takes enormous computing power), they can die, they show emotions, talk, etc. In one of the theaters a million avatars can visit rock or graphics concerts.

There are also agents living in this world. For example, *The Black Sun*, a popular bar for which membership is needed. It has bartenders and bouncer agents. Companies have receptionist agents, usher ladies and security agents. Obviously, every visitor of the company will see a 'matching' agent (e.g., a Japanese business man will see a Japanese looking receptionist). Rich people can afford themselves very expensive agents. What to think of a perfectly rendered geisha agent that can rub the back and shoulders of your alter ego avatar?

In this paper we surveyed the developments in the design of virtual interest communities, the way people can represent themselves in these communities and how they can explore and interact, not only with each other, but also with community agents with task and domain knowledge. The developments we described and that are in line with the situations described by Stephenson were illustrated with examples from our research on a virtual music theatre that can become a meeting place for people interested in theatre in general.

References

[1] M. Bertolo, P. Maninetti & D. Marini. Baroque dance animation with virtual dancers. *Eurographics '99*, Short Papers and Demos, Milan, 1999, 117-120.

[2] B. Damer. Avatars!: Exploring and building virtual worlds on the internet. Peachpit Press, Berkeley, 1998.

[3] S. Das & D. Grecu. COGENT: Cognitive agent to amplify human perception and cognition. 4th *International Conference on Autonomous Agents 2000*, C. Sierra et al. (eds.), Barcelona, 2000, 443-450.

[4] M.D. Dickey. 3D Virtual Worlds and Learning: An analysis of the impact of design affordances and limitations in Active Worlds, Blaxxun Interactive, and OnLive!Traveller. Ph.D. Thesis, Ohio University, 1999.

[5] D. Lie, J. Hulstijn, R. op den Akker & A. Nijholt. A Transformational Approach to NL Understanding in Dialogue Systems. Proc. *NLP and Industrial Applications*, Moncton, 1998, 163-168.

[6] B. Laurel. *Computers as Theatre*. Addison-Wesley 1991; 2nd edition 1993.

[7] J.-S. Monzani & D. Thalmann.. Verbal communication: using approximate sound propagation to design an inter-agents communication language. 4th *Int. Conf. on Autonomous Agents 2000*, C. Sierra et al. (eds.), Barcelona, 2000, 235-236.

[8] S.R. Musse & D. Thalmann. From one virtual actor to virtual crowds: requirements and constraints. 4th *Int. Conf. on Autonomous Agents 2000*, C. Sierra et al. (eds.), Barcelona, 2000, 52-53.

[9] A. Nijholt & J. Hulstijn. Multimodal Interactions with Agents in Virtual Worlds. In: *Future Directions for Intelligent Information Systems and Information Science*, N. Kasabov (ed.), Physica-Verlag: Studies in Fuzziness and Soft Computing, 2000.

[10] A. Nijholt & H. Hondorp. Towards communicating agents and avatars in virtual worlds. In: Proceedings *EUROGRAPHICS 2000*, A. de Sousa & J.C. Torres (eds.), Interlaken, August 2000.

[11] R. Price, C. Douther & M.A. Jack. An investigation of the effectiveness of choreography for the portrayal of mood in virtual environments. 4th *International Conference on Autonomous Agents 2000*, C. Sierra et al. (eds.), Barcelona, 2000, 54-55.

[12] M. Reany. The Theatre of Virtual Reality: Designing Scenery in an Imaginary World. *Theatre Design and Technology*, Vol. XXIX, No.2, 1992, pp. 29-32.

[13] C. Reeve. Presence in Virtual Theatre. *BT Presence Workshop*, BT Labs, 10-11 June 1998.

[14] G. Reitmayr et al. Deep Matrix: An open technology based virtual environment system. *The Visual Computer Journal* 15: 395-412, 1999.

[15] D. Rousseau & B. Moulin. Mixed initiative in interactions between software agents. Proc. 1997 Spring Symposium on *Computer Models for Mixed Initiative Interaction*, AAAI Press, Menlo Park, CA, March 1997

[16] N.M. Sgouros. Involving the audience in distributed, interactive entertainment systems. In: *Multimedia '99*. Proc. of *Eurographics* Workshop. N. Correia et al. (eds.), Springer, 2000, 113-120.

[17] R. Vertegaal et al. Why conversational agents should catch the eye. Proc. *CHI 2000*, 2000, 257-258.