M.F. Sebaaly (Ed.), Proceedings of the International NAISO Congress on Information Science Innovations. Dubai, March 17-21, 2001. ICSC Academic Press, Canada/The Netherlands, 347-353.

Agents for Navigating Virtual Reality E-Commerce Environments[1]

Anton Nijholt and Betsy van Dijk
University of Twente, PO Box 217
7500 AE Enschede, the Netherlands
anijholt@cs.utwente.nl

Abstract: *We report about work in progress on capabilities of navigation agents in virtual environments. These agents may help a user to explore the environment: where to go, where to buy, where to find, etc. The agents may make suggestions based on a user profile, earlier visits and earlier given advice and these agents may exchange information with other types of agents that may be useful for the visitor of the environment or for the owner of the environment. In our case the environment consists of a multi-purpose theatre building where visitors can get information about performances, but also can make reservations for these performances (buy tickets). Environments like these can be viewed from different perspectives, e.g., an E-commerce perspective (an electronic shop), a community perspective (a market and performance place) where people can meet each other, or an edutainment perspective (people can learn about professional activities that take place, e.g., learning to play an instrument, take part in a performance, design a choreography, etc.). Here we concentrate on navigation in these virtual environments. How can we assist the user in finding its way in such complex and information rich environments? Presently we maintain different subprojects in which navigation aspects of virtual environments are studied. Common characteristics of these subprojects are the use of embodied conversational agents and multimodal interaction between users and embodied agents.*

1 Introduction

We discuss how we can support users to navigate in virtual environments and to assist them in arriving at the information they are interested in, at the products they want to buy, at the people they want to meet and at the demonstrations they want to see. As a laboratory for implementing and integrating ideas we are using a virtual reality theatre environment that has been build a few years ago in our research group. The environment has been built using VRML (Virtual Reality Modeling Language) and is accessible on WWW. In this environment agents have been defined, allowing (primitive) dialogues with users or visitors of the system. The original system has been evolved by adding more facilities 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. The current paper is an example of our growing interest in making environments like these accessible for larger audiences.

Visitors can explore our 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. Karin employs a dialogue system that gets its information from a database filled with all performances that will take place during one year in the theatres of Enschede, our home town. An environment like this can be looked at from different viewpoints. In Nijholt [14] we reported about the e-commerce aspects of this environment, in Nijholt [15] about the community building aspects. Here we concentrate on the navigation in these virtual environments. Presently we maintain four different subprojects in which navigations aspects are studied. In the near future we expect to be able to integrate these four approaches in one or two prototype navigation agents that will be used in practice.

2 Agents in Virtual Environments

In our environment we can have different human-like agents. Some of them are represented as communicative humanoids, more or less naturally visualized avatars

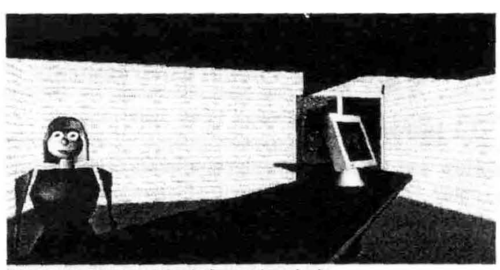

Figure 1: Karin at the information desk

[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.

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.

We envision a future environment where vistors can get their information, perform transactions and find other people and agents they can communicate with and that maybe can assist them in doing whatever they want to do in virtual or real life. It seems natural that in such an environment the representation of other visitors and of domain agents are human-like and are animated. We should be able to identify different visitors and agents in order to address them for particular questions and we want to recognize our friends, colleagues and our favorite salesperson. When we communicate we expect feedback that helps to regulate the conversation. Hence, in addition to speech recognition and language understanding, non-verbal interaction behavior, which requires modelling of facial expressions, eye movements and hand gestures, is needed. In situations where an agent is expected to demonstrate a product or to teach us a task which involves moving and manipulating virtual objects then we should be able to see how this agent uses his arms, hand and legs in order to be able to imitate it.

Our environment is a reasonable realistic representation of an actual theatre. It has been built according to the design drawings of the architects of the building. That means we have different performance halls, information stands, lounges, a second and a first floor, lounges and coffee stands, different entrances, etc. There are posters and information boards that are updated (partly automatically because when time passes the database can be accessed to get new information to be displayed on poster boards) and, since we have a multi-user shell avalaible, visitors can not only meet domain-provided artificial agents but also other (embodied) visitors with whom they can talk using a chat window.

3 Navigation in Virtual Environments

3.1 Introduction

How can we look at navigation in an environment like ours? From a general point of view navigation takes place in information spaces, where these spaces may range from data bases, hypertext and hypermedia systems, immersive and non-immersive virtual environments to the real world. Navigation has similarities to other problem solving tasks, but generally it is defined in such a way that some kind of locomotion in space is involved. For problem solving we can nevertheless say that we have to find our way to a solution given a space of constraints.

As mentioned in Höök et al. [6], in order to make navigation support available to a user it is necessary to view an information space among its many dimensions: the tasks that have to be performed, size and stability, moderated or non-moderated, the presence of other users, Euclidean properties, the richness of presentation, etc. In addition it is necessary to look at individual differences of the users, e.g., their different cognitive (spatial) abilities.

A generally accepted definition of navigation involves the following parts (Downs and Stea [4]):

- orienting oneself in the environment,
- choosing the correct route,
- montoring this route, and
- recognising that the destination has been reached.

Navigation originally was defined as the process of moving through an environment. In Darken & Sibert [2] this definition of navigation is reviewed to include the process of wayfinding (determining a path to be traveled). They extend the definition even more with a reference to the aids and cues people need for successful navigation in virtual environments, to arrive at the following definition: "Navigation is the process by which people control their movement using environmental cues and artificial aids such as maps so that they can achieve their goals without getting lost" (Darken & Sibert [2]).

Other authors have made additions to the four parts of the definition mentioned above. The main observation is that these parts are too goal-directed. Whereas the navigation activities in traditional information spaces mainly comprise searching for information, especially with the development of hypermedia and virtual worlds the attention for exploration increases. The user or visitor may want to explore a virtual space, may want to learn about it, may want to meet other people or may want to identify objects or events in a space and therefore be involved in activities that do not always require choosing or monitoring the 'correct' route or require correct orientation at each moment in time. This exploration may induce pleasure and pleasure may be part of the goal rather than aiming only at being efficiently guided to a destination. A similar observation can be made when the environment is also meant to learn about objects or to bring us in contact with other people (who may be there or who may have been there). Obviously, goal-oriented search and exploration may become connected during navigation: exploration may invoke goal-oriented search activities, but the opposite direction is conceivable as well. Darken & Sibert [3] present a classification of wayfinding tasks in which they distinguish exploration (defined as a wayfinding task without a target) from search activities, where search activities are subdivided in naive search (the navigator has no a priori knowledge of the whereabouts of the target) and primed search (the navigator knows the location of the target).

3.2 Navigation in Virtual Worlds

It has often been noted that navigation in virtual environments is difficult. Unfamiliar environments are always difficult to navigate, whether they are virtual or real (Vinson [19]), but in virtual environments more problems may be expected. Problems associated with wayfinding in virtual worlds occur because of several reasons. First of all, because in virtual environments in general less sensory (visual, auditory, locomotive) detail is presented, they contain fewer spatial and locomotive cues than real environments. Furthermore, the concept of navigation in hyperspace may have completely different physics from navigation in the physical world. For instance node-link representations in hypermedia and on the web permit discrete movement while movement in Euclidean space is continuous (Spence [17]. Input by keyboard and mouse allows the user to move and to rotate, to jump from one location to another, to interact with objects and to trigger them. The user also has the ability to view the world from different perspectives (Satalich [16]).

Visitors of a virtual world often will encounter situations where they will not be able to rely on navigation skills acquired in the real world, as familiar kinds of cues are inefficient or inappropriate. Navigation problems like disorientation, loss of overview, difficulty to return to a location visited before or to refind an object found before, can lead to dissatisfaction, frustration and eventually discontinued use of that environment (Nash et al. [12]). Navigational support is needed and many researchers nowadays study navigation and wayfinding in the hope their findings will be applicable to ease navigation in virtual environments. The process of determining a path (wayfinding) is inherently cognitive in nature (Nash et al. [12]). Many studies of navigation focus on understanding the knowledge and abilities it requires (e.g. Krieg-Brückner et al. [8]; Werner et al. [21]; Spence [17]). These studies often compare real world wayfinding to wayfinding in virtual worlds. Theories on spatial knowledge and navigational awareness seem to be relevant for wayfinding in both the physical world and virtual worlds. In spatial knowledge theory two types of knowledge are distinguished:

- survey knowledge is characterized as the ability to conceptualize the space as a whole. This knowledge is map-like in nature: many routes and landmarks are combined into a cognitive map of the environment.
- procedural or route knowledge is defined as the sequence of actions required to follow a particular route. A person who has only procedural knowledge of a route can go from one landmark to another on that route but does not recognize alternative routes.

In the literature it is agreed upon that survey knowledge is the key to successful effective navigation (e.g. Darken & Sibert [3]; Nash et al. [12]). In order to be able to create a cognitive map of the virtual environment the user should be able to orient him/herself in space and build up landmarks, route and survey knowledge (Volbracht & Domik [20]). Survey knowledge can best be attained by exploration. Part of the knowledge can be attained through map or picture study alone but for a correct orientation and location of landmarks personal experience through exploration is necessary. Navigators with complete survey knowledge are said to possess navigational awareness.

Based on the role of spatial knowledge in real world wayfinding tasks, Darken & Sibert [3] assert that real world wayfinding principles can be succesfully applied to construct aids for wayfinding in virtual worlds. They present a set of design principles for wayfinding augmentations to virtual worlds. These design principles partly concern the organisation of the environment (divide large-scale worlds into distinct parts that are simply organized), partly they propose the addition of map-like information to the environment. They advise to show paths, landmarks, and the user's position. Furthermore they advise to orient the map with respect to the user such that "the forward-up equivalence principle is accomodated".

3.3 Resembling the Real World

Being in a multi-user environment where we can meet other visitors or where others have left cues to reach an interesting destination allows direct, respectively indirect, navigation. We can ask people where to go and to find a location (or certain information) or we can find routes taken by other people with maybe an interest profile similar to ours. As mentioned by Höök et al. [6], navigation "places the user inside the system", rather than having the user concerned with the direct manipulation of opjects (opening and closing files, e.g.) on the screen. Obviously, this is especially true in rich visualised environments, as is usually the case in (sometimes geographic properties preserving) 3D and virtual reality environments.

As mentioned, goals can be more or less definite. This shows in the real world when buying the usual pack of milk, a news paper, a pair of shoes, a toy, ingredients for tonight's dinner, etc. In the latter case, during shopping and depending upon what is available, what special offers there are, what time is available for shopping, etc., the decisions are made and the dinner is composed. Hence, apart from navigating in a physical space (the actual mall, the shop, the web pages shown on the screen), there is navigation in a space with product information and a space containing constraints that need to be considered in order to make decisions. In information-rich virtual reality environments will continuously have to make similar decisions as in real-world situations. If, in addition, these environments are inhabited by synthetic help agents, agents with self-interest and agents 'just' presenting other users, then it is certainly true that navigation takes place in a world that in many aspects goes beyound the geographic and social space of the real world.

In the next section we describe some of our projects in which navigation and navigation assistance is studied. These projects already apply several of the design principles for navigation aid mentioned in this section while other navigation aids will be added and tested in the near future. Furthermore these projects have in common that the user is assisted by a personal agent. This advisor agent or navigation agent can give advice or make suggestions about interesting places to go to, how to get there, where to find things or how to solve certain problems. The user can be guided (route guidance/guidance in problem solving) or s/he can be transported directly, in which case the user hands over control to the agent.

4 Navigation and Assistance Agents

Presently four approaches are followed in our research on navigation aids in virtual environments. These approaches, unfortunately, have to be followed in different projects. The first one we describe is concerned with offering the user speech and language interaction with the virtual environment. The second project is the U-WISH (Usability of Web-based Information Services for Hypermedia) project in which we participate as members of the Dutch Telematics Institute. The third project and fourth project are different in the sense that they do not explicitly deal with navigation. However, we think that the topics that are dealt with in these projects are that closely related to the navigation issues we are interested in, that they should be mentioned shortly in this overview. Moreover, especially in these two projects the role of embodied agents becomes visible. One is the Jacob project which we do as members of the VR-Valley Foundation, an initiative which aims at establishing a regional knowledge center on virtual reality in the Netherlands. The project is concerned with problem solving assistance in a virtual environment. A final approach is done in cooperation with Queens University in Kingston, Canada. In that project we use eyetrackers to experiment with gaze detection in virtual environments. We hope to be able to combine the results of the different approaches in a future design of navigation agents in our virtual environments.

4.1 Navigation Using Speech and Language

Already from the start it turned out that some non-professional users had problems navigating in our virtual environment. Therefore we introduced a navigation agent that could be addressed in limited natural language using the keyboard or spoken utterances. Apart from the well-known shortcomings of state of the art speech technology it turned out to be a useful addition. In the system it is left to the user to choose between interaction modes (speech, keyboard, mouse) or to use all, sequentially or simultaneously.

In general, a smooth integration of the pointing devices and speech in a virtual environment requires that the system has to resolve deictic references that occur in the interaction. Moreover, the navigation agent should be able to reason (in a modest way) about the geometry of the world in which it moves. The navigation agent knows about the user's coordinates in the virtual world and it has knowledge of the coordinates of a number of objects and locations. This knowledge is necessary when a visitor refers to an object close to the navigation agent in order to have a starting point for a walk in the theatre and when the visitor specifies an object or location as the goal of a route in the environment. The navigation agent is able to determine its position with respect to nearby objects and locations and can compute a walk from this position to a position with coordinates close to the goal of the walk.

Verbal navigation requires that names have to be associated with different parts of the building, objects and agents. Users may use different words to designate them, including references that have to be resolved in a reasoning process. The current agent is able to understand command-like speech or keyboard input. Otherwise it hardly knows how to communicate with a visitor. The phrases to be recognized must contain an action (go to, tell me) and a target (information desk, synthesizer). It tries to recognize the name of a location in the visitor's utterance. When the recognition is successful, the agent guides the visitor to this location. When the visitor's utterance is about performances the navigation agent makes an attempt to contact Karin, the information and transaction agent (Lie and Leijdekkers [9]). Because of ownership problems of the commercial software that is used (Speech Pearl, Philips) the navigation agent has not been included in the publicly accessible websites that have been made available for our system. Moreover, we felt that having to deal with the limited properties of the current level of speech technology restricted us too much in developing more general ideas about navigation support.

More recently a much more advanced natural language accessible navigation agent has been built where the natural language input has to be typed by the user (van Luin [10]) and where natural language references can be made to objects and locations visible on a floormap or pointed at by the mouse. The navigation agent knows about (or is able to compute), among others:

- Current position and what is in the eyesight of the user;
- Objects and the properties they have; geometric relations between objects and locations;
- Possible walks towards objects and locations;
- Some knowledge of previously visited locations or routes;
- The action it is performing (or has performed);
- Some knowledge of the previous communication with the visitor.

Two representative examples of dialogues with the agent are shown below:

Dialogue 1:

Visitor: Can you bring me to the information desk?
Agent: I have marked the position on the map.
Now I will bring you there.

Dialogue 2:

Visitor: [Clicks on an object on the map]
What is this?
Agent: That is an exposition.
Visitor: Where is it?
Agent: You can find it in the lounge.
Visitor: Let's go there.
Agent: I bring you there.

Whenever the agent has computed a route from the current position of the visitor to a desired object or location then it sends this route to the VR browser agent which then guides the visitor's viewpoint along this route to the destination.

4.2 Navigation in the U-WISH Project

In the U-WISH project (Neerinck et al. [13]) cognitive engineering techniques are used to develop and test support concepts for networked user interfaces and to derive HCI guidelines based on the test results. One of the test services being used in the U-WISH project is the virtual music center. In the context of this project a new agent-based navigation assistant has been built. Rather than exploring the problems associated with addressing such an agent using speech and language, here the emphasis is on the possibility to obtain an evaluation framework in which different kinds of user interfaces can be compared. This required some simplifications on our side, but also some useful extensions, e.g. user profiles.

It is clear that in many situations we can expect different user interaction behavior and different user preferences with respect to the 'content' that is offered. These differences follow from different interests, background, culture, intelligence and interaction capabilities of users. These issues can become part of a user profile (obtained by learning, by assuming or by asking), help the system to anticipate the users preferences and even help to guide a user's avatar acting in the virtual environment. For experimental purposes the user profiles in the U-WISH project are fixed. They just contain a few fields containing, among others, name, profession and interests of a user.

In this project, in the user's browser we have an 'eavesdropper' that listens to the interactions of the user with the virtual environment (our virtual music center) and sends them to the server. For each user the server has an administrator agent that creates (or loads) a user profile, an event history and an advice history. Moreover, it creates a number of sub-agents. Events coming from the client are received by the administrator agent, entered into the event history and then send to an appropriate sub-agent. Responses from a sub-agent are logged in the advice history and send to the client's virtual music center. For instance, there is a sub-agent called the PositionAgent, which generates responses based on the position (triggered when the user passes a sensor in the virtual environment), the event history and the profile of a user. Similarly, there is a sub-agent called the DialogAgent, which monitors the dialogue with Karin for certain keywords. The responses by these and other possible sub-agents take the form of suggestions to the user, which, at this moment, are displayed,

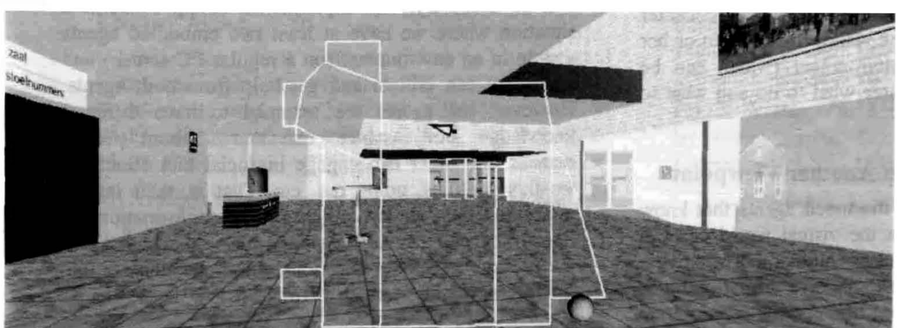

Figure 2: Navigation floor map on top of VR world with arrow displaying position and viewpoint

in an advice window. This window may contain text, hyperlinks and internal links to other parts of the virtual environment. The current agents are rule-based, but as long as they comply with the input/output conventions in their communication with the administrator agent more sophisticated agents can be introduced.

During the U-WISH navigation experiments that are now in preparation tasks have to be performed. They are embedded in scenarios about fictive users. Some of the tasks are open (find some general information within a certain limit of time), others are closed (find a specific piece of information). In the current experiments we are not yet using the natural language accessible navigation agent we discussed in section 4.1. As mentioned, the aim for this moment is to obtain an evaluation framework in which different interfaces (employing different agents), can be compared. Nevertheless, some 'passive' support is provided by having our theatre's floormap imposed upon the virtual world as shown in Figure 2. The moving arrow denotes the current position of the user and shows in which direction in the 3D world s/he is looking at this moment.

4.3 Navigation in the Jacob Project

In the Jacob project (cf. Evers and Nijholt [5]) we have the task to design an animated agent, which is called Jacob, in virtual reality, which gives instruction to the user. In this project software engineering plays a prominent role. We apply object-oriented techniques, design patterns and software architecture knowledge. In the architecture we have separated the concerns of the 3D visualization from the basic functionality, which follows from a task model, an instruction model and a user model. Presently, the task and instruction model form Jacob's mind, a control system that observes the world and tries to reach specific objectives by having Jacob perform a certain task (e.g., show the user what to do next) or to produce an utterance to direct the user. Presently Jacob's task is to teach a user to solve the Towers of Hanoi problem (cf. Figure 3). This is chosen as an example task since we think that the design solutions found there can be generalized very well and applied to the situation where some embodied agent like Jacob will be integrated in the virtual music center where it can interact with the visitor to help him or her to navigate through the environment (which can be considered as to teach the user what to do and what to find where).

4.4 Attentive Agents: Yet Another Viewpoint

In sections 4.1 and 4.2 we discussed agents that know about the user's position in the virtual world or even know in which direction the user is looking or what is in his or her eyesight. For non-immersive virtual reality it is useful to make a distinction between the user looking at a PC screen seeing a virtual world and being able to move his gaze over the screen without changing the view of the virtual world that is presented on the screen and the user presented in the virtual world by a virtual camera at a certain position (e.g. the mouse pointer coordinates) and looking in a certain direction in the virtual world. This is the difference between the first peron perspective and the third person perspective. Although there is chance of becoming confusing, one

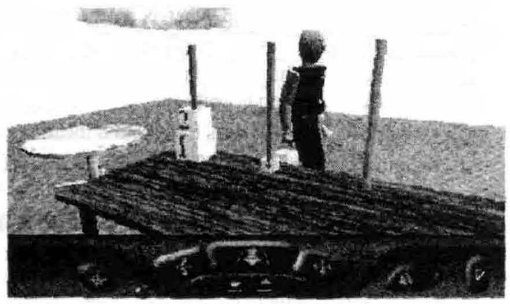

Figure 3: Jacob shows the user what to do next

might think of navigating in a virtual world by using gaze detection. Where is the user looking at? Does that mean he wants to go to that place? Should a navigation or assistance agent start explaining about that object? Especially when we move beyond the current desktop computer and instead look at attentive environments where joint voice and gaze information can (unambiguously) activate one or more devices and agents (see Matlock et al. [11]) we enter a situation where we both have viewpoints in a virtual environment and an additional gaze guided steering mechanism from a third person perspective.

Several of the projects we are currently engaged in concern the use of gaze detection and what to conclude from fixing the gaze during a certain amount of time and changing the gaze direction because of activities on the screen. A computer system can establish where the user is looking at by means of a desk-mounted eye or head tracking system. When the screen is big enough head tracking may be sufficient, otherwise eyetracking will be necessary. Our experimental set-up concerns a situation where we have at least two embodied agents visible in an environment (on a regular PC screen) and where we can talk to and get help from both agents. However, the agents are assumed to have different knowledge. For instance, one knows about performances, the other can engage in social talk about the weather, politics, going out, etc. That is, their intelligence and their ability to give sensible information are domain-dependent. As long as is clear to which agent the user addresses her questions, these questions can be interpreted in the domain maintained by the particular agent. When the user is able to recognize where the agent stands for s/he will addres the agents with questions concerning that particular domain, making it more easy to interpret the user's utterances. We are currently working to implement our earlier findings on gaze behavior (Vertegaal et al [18]). Using this technology the system is capable of detecting where the user is looking (a particular agent, a particular object . .) and may use this information to understand a user's next or previous utterance better (it can disambiguate ambiguous utterances) and to ask relevant questions or to make relevant suggestions.

5 Conclusions

We surveyed our approaches to the design and implementation of navigation agent in virtual environments for information services and transactions. The approaches take place in different subprojects each emphasizing a different aspect of a navigation agent. The approaches will be integrated in the design and implementation of some agents in the course of 2001. The integration will be influenced by experiments we perform with real users in a usability laboratory and with users that will be taken from our own student population.

Acknowledgements: We gratefully acknowledge help from J. Zwiers, M. Evers, R. op den Akker, H. Hondorp, J. van Luin and C. Eyckelhof.

References

[1] D. Benyon & K. Höök. Navigation in information spaces: Supporting the individual. S. Howard, J. Hammond & G. Lindegaard (eds.), *Human-Computer Interaction: INTERACT'97*, Chapman & Hall, 1997, 39-46.

[2] R.P. Darken & J.L. Sibert. A toolset for navigation in virtual environments. In Proceedings of *UIST '93*, Atlanta GA, 1993, 157-165.

[3] R.P. Darken & J.L. Sibert. Wayfinding strategies and behaviors in large virtual worlds. In Proceedings of *CHI '96*, 1996, 142-149.

[4] R. Down & D. Stea. Cognitive Representations. In: R. Downs & D. Stea (eds.), *Image and Environment*, Chicago: Aldine, 79-86.

[5] M. Evers & A. Nijholt. Jacob - an animated instruction agent for virtual reality. In: *Advances in Multimodal Interfaces - ICMI 2000*, Third Int. Conf. on Multimodal Interfaces, Beijing, China, October 2000, LNCS 1948, T. Tan, Y. Shi & W. Gao (eds.), 526-533.

[6] K. Höök et al. Towards a Framework for Design and Evaluation of Navigation in Electronic Spaces. *Persona* Deliverable for the EC, 1998.

[7] A. Kendon. Some functions of Gaze Direction in Social Interaction. *Acta Psychologica* 32, 1-25.

[8] B. Krieg-Brückner et al. A taxonomy of spatial knowledge for navigation and its application to the bremen autonomous wheelchair. C. Freksa et al. (eds.), *Spatial Cognition*, LNAI 1404, 1998, 373-397.

[9] D. Lie & B. Leijdekkers. Navigatie-agent voor het Virtuele Muziekcentrum (VMC). TKI Internal Report, University of Twente, 1998.

[10] J. van Luin. Navigatie in het Virtuele Muziekcentrum. M.Sc. Thesis, December 2000.

[11] T. Matlock et al. On gaze and speech in attentive environments. *Advances in Multimodal Interfaces*, Third Int. Conf. on Multimodal Interfaces, Beijing, China, 2000, LNCS 1948, T. Tan et al. (Eds.), 1-7.

[12] E.B. Nash et al. A review of presence and performance in virtual environments. *International Journal of Human-ComputerInteraction*, 12:1-41, 2000.

[13] M.A. Neerincx et al. U-WISH; Web usability: methods, guidelines and support interfaces. TNO-report TM-99-D005, TNO Human Factors Research Institute.

[14] A. Nijholt. Virtual malls for commerce and entertainment. In: *Trends in Distributed Systems: Towards a Universal Service Market.* Third Int. IFIP/GI Working Conference, September 2000. LNCS 1890, C. Linnhoff-Popien & H.-G. Hegering (eds.), 326-333.

[15] A. Nijholt. Towards virtual communities on the Web: Actors and audience. ICSC Symp. on Interactive and Collaborative Computing (ICC'2000). University of Wollongong, Australia, December 2000, to appear.

[16] G. Satalich. Navigation and wayfinding in virtual reality. Master's thesis, University of Washington, 1995.

[17] R. Spence. A framework for navigation. *International Journal of Human--Computer Studies*, 51:919-945, 1999.

[18] R. Vertegaal, R. Slagter, G. van der Veer and A. Nijholt (2000). Why conversational agents should catch the eye. Proceedings *CHI 2000*, 257-258, 2000.

[19] N. Vinson. Design guidelines for landmarks to support navigation in virtual environments. In Proceedings *CHI '99*, 278-285, 1999.

[20] V. Volbracht & G. Domik. Developing effective navigation techniques in virtual 3d environments. J.D. Mulder & R. van Liere (eds.). *Virtual Environments 2000.* Eurographics Workshop, 2000, Springer, 55-64.

[21] S. Werner et al. Modelling navigational knowledge by route graphs. C. Freksa et al. (eds.), *Spatial Cognition II*, LNAI 1849, Springer, 2000, 295-316.