Jacob: A Web-based Learning Environment Using Virtual Reality

Anton Nijholt
Centre of Telematics and Information Technology
University of Twente, PO Box 217
7500 AE Enschede, the Netherlands
anijholt@cs.utwente.nl

Marc Evers
Software Engineering Research Centre
PO Box 424, 3500 AK Utrecht,
the Netherlands
evers@serc.nl

Abstract

This paper gives an overview of the Jacob project. This project involves the construction of a 3D virtual environment where an animated human-like agent called Jacob gives instruction to the user. The project investigates virtual reality techniques and focuses on three issues: the software engineering aspects of building a virtual reality system, the integration of natural language interaction and other interaction modalities, and the use of agent technology. The Jacob agent complies with the H-Anim standard. It has been given a task model and an instruction model in order to teach the user a particular task. The results of the project can be generalised so that the agent can be used to instruct other tasks in other virtual environments. One of the applications we foresee is the use of a navigation agent which advises the user in a VR environment or, using a PDA-like device, about a real environment, thatcan be perceived in reality and in a matching virtual reality.

1. Introduction

The Jacob project investigates the application of virtual reality techniques and involves the design and construction of an animated agent in a 3-dimensional virtual environment. The agent is called Jacob and provides instruction and assistance for tasks that the user has to learn to perform in a virtual environment. The interaction between the user and Jacob is multimodal: interaction takes places through natural language as well as by performing actions. The use of a lifelike agent in an interactive learning environment has a strong positive impact on students, which is shown by an empirical study performed by Lester et al. Such an agent can increase both the learning performance and the student's motivation [2].

The Jacob project involves an integration of knowledge from different disciplines, like intelligent tutoring systems, virtual reality, intelligent agent technology, natural language processing, and visualisation and animation techniques. It is a pilot project of the VR-Valley Twente Foundation, which aims at establishing a regional knowledge centre on virtual reality in the Netherlands [6].

The Jacob system is a web-based learning system that has been developed for the Jacob project. In this system, the Jacob agent teaches the user the Towers of Hanoi game, a classic toy problem from the field of artificial intelligence. In this game, the user has to move a stack of blocks from one peg to another using an auxiliary peg. The user can move single blocks from peg to peg; it is not allowed to place a larger block on top of a smaller one. In the Jacob system, the user and Jacob can communicate through manipulation of blocks and through written language utterances. Jacob also shows emotions, e.g. satisfaction or anger. In Figure 1 we see Jacob in action in the virtual world.

The Jacob system has been written using VRML 2.0 and Java. It runs as an applet in a web browser. The system forms a basis for learning over internet (or intranet) using virtual reality. The Jacob system is currently single-user, but a multi-user version has been

Figure 1: The Jacob Agent Moving Blocks

planned. In this multi-user version, several users populate the same virtual environment and collectively perform tasks, assisted by Jacob.

2. Software Architecture of the Jacob System

Figure 2 gives an overview of the software architecture of the Jacob system. This architecture is described in detail in [1]. We will give a short description of the different components.

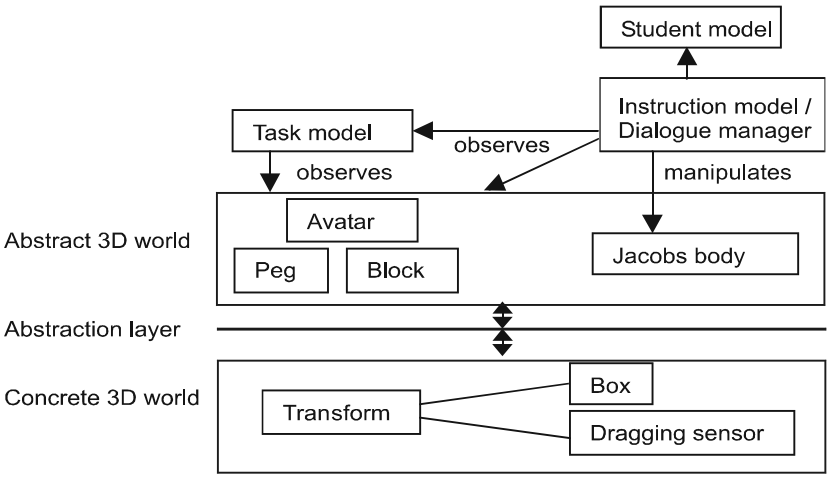

Figure 2: Architecture of the Jacob System

- The *Abstract 3D World* is a high level model of the virtual environment. This model contains semantic knowledge of the world: it contains representations of the user (avatar), Jacob, and the different objects like blocks and pegs. This component is responsible for gravity simulation and collision avoidance.
- The *Concrete 3D World* takes care of the visualisation of the virtual environment. It manages the different VRML object structures and encapsulates the interface to the VRML browser. This component contains very little semantics: objects are defined here in terms of geometry, colours, textures, etc.
- The *Task Model* encapsulates knowledge about the objects involved in the task and how the task should be performed. For the Towers of Hanoi task, it contains rules that describe how the Towers of Hanoi task has to be solved. The task model observes the abstract 3D world and interprets the state of the objects in the world according to the rules of the task. It also defines task related constraints. For example, the Hanoi task defines the constraint that larger blocks should not be placed on top of smaller ones.
- The *Instruction Model / Dialogue Manager* encapsulates instruction knowledge and natural language understanding capabilities. It observes the world, the task and the user's actions and utterances. It can manipulate the world and the user through utterances and a limited set of actions. These actions include shaking Jacob's head, performing the next step of the task, telling what the next step is, and letting Jacob pick up the next object and hand it to the user. We have found that instruction knowledge and natural language dialogue are tightly connected: interaction in the Jacob system is a mix of actions and utterances, performed by both the user and the Jacob agent.
- The *Student Model* is a model of the user's characteristics and performance. It has not yet been implemented in the Jacob system.

Because of the complexity of the dialogue system and the size of the grammar and lexicon, Jacob's dialogue system cannot be executed as part of the applet on the user's machine and runs on a separate server machine. There is a socket connection between the dialogue system and the Jacob applet by which events and actions are exchanged.

3. Future Work

We have identified a number of directions for future work. Extending the Jacob system to a multi-user system is an important direction. In this way, the Jacob system will evolve into a distributed, collaborative learning environment.

To determine the lexicon and grammar for the user utterances, we will perform Wizard of Oz experiments. In a Wizard of Oz experiment, users work with the Jacob system in an experimental setting where Jacob and Jacob's responses are controlled by a human. The dialogues are logged and analysed. We have already constructed a distributed version of the Jacob system where Jacob can be controlled from a remote system. Wizard of Oz experiments can also be useful to find out what kind of actions and nonverbal behaviour to expect from users.

We will extend and generalise the Jacob system to other tasks, as the Towers of Hanoi game has a number of restrictions. For this purpose, we are investigating how to integrate Jacob in the Virtual Music Centre [3]. There, Jacob can give instruction for a task like navigating through the virtual music centre (to teach the user what to do and to find where) or to make suggestions where to go [7]. Once we have developed such an agent, it can also be used to assist the user to navigate in the real world. For this purpose an

Figure 3: Pocket PC Navigation Assistance

interesting development is the use of virtual reality techniques on PDAs and cell phones. Recently, a VRML browser for PDAs has been released, Pocket Cortona from ParallelGraphics [4], see Figure 3. A second application of Jacob that we would like to investigate is a tutoring system where object-oriented programming concepts like objects and classes have a three-dimensional representation in virtual reality. A program or a design is represented by a spatial structure that can be manipulated by the user. In this system, Jacob instructs the user in object-oriented programming and design techniques.

The Jacob system can be extended to handle other interaction modalities like gestures and gaze detection [5]; we are thinking of using a dataglove, which allows the user to manipulate objects in the virtual world more directly and which can be used to point at objects, and an eye tracker. Speech recognition and speech synthesis can also be added to the system. We want to provide more natural animation for Jacob, including a more elaborate visualisation of emotions; we will investigate impact of this on human computer interaction.

References

[1] Evers, M., Nijholt, A., "Jacob, an agent for instruction in virtual reality", submitted to *Themes in Education*.

[2] Lester, J.C., Converse, S.A., Kahler, S.E., Barlow, S.T., Stone, B.A., Bhogal, R.S.: "The persona effect: affective impact of animated pedagogical agents". In: *CHI'97 Proceedings* (1997) 359-366

[3] Nijholt, A. and Hulstijn, J., Multimodal Interactions with Agents in Virtual Worlds, In N. Kasabov (Ed.), *Future Directions for Intelligent Information Systems and Information Science*, Physica-Verlag: Studies in Fuzziness and Soft Computing (2000), 148-173.

[4]. ParallelGraphics, *Pocket Cortona*, http://www.parallelgraphics.com/products/cortonace

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

[6] VR-Valley Twente homepage, http://www.vr-valley.com/

[7] Zwiers, J., van Dijk, B., Nijholt, A. and op den Akker, R. Design issues for navigation and assistance in virtual environments. In Proceedings *Learning to Behave*, University of Twente, October 2000, 119-132.