

Tele-Experiments – Experiments on Spatial Cognition using VRML-based Multimedia

Galia Givaty *

Hendrik A.H.C. van Veen †

Chris Christou ‡

Heinrich H. Bülthoff §

Max-Planck-Institute for Biological Cybernetics
Spemannstraße 38, 72076 Tübingen, Germany



Abstract

This paper describes an application of VRML-based multimedia to the domain of spatial cognition experiments. Typically, such experiments involve human subjects that are presented with graphical stimuli in a laboratory setup. Larger numbers of subjects and stimuli give more weight to the outcomes of the experiment, but in practice the number of subjects is limited by practicalities like time constraints. This paper describes a way to overcome this limitation: use the internet to bring the experiment to the subject instead of the other way around. We are developing a web page with psychophysical experiments, using VRML as a means for presenting three-dimensional graphics to simulate the conditions of a laboratory experiment. While using the web for psychological experiments is not a novel idea, combining that with VRML is new and extends the possibilities for experiments. In this paper we outline the power of using the web and describe an experiment using VRML, highlighting technical aspects of the problems encountered and the solutions taken to overcome them.

CR Categories: I.3.7 [Three-Dimensional Graphics and Realism]: Virtual reality; J.4 [SOCIAL AND BEHAVIORAL SCIENCES]: Psychology

Keywords: VRML, Internet, Tele-Experiments, Psychology, Spatial Cognition, Virtual Reality

1 INTRODUCTION

One approach towards a better understanding of how our brains work involves doing perceptual, cognitive and behavioral studies with human subjects. Psychologists investigate human cognition by making inferences from subjects' responses to particular stimuli. The main problem with this approach is the control of unwanted cues or variables (see [6]). In a typical experiment, groups of subjects are asked to respond in a strictly specific manner to a stimulus containing only the pertinent cue of the experiment. The pattern of responses reveals information about the underlying neural mecha-

nisms used to perform the task. When inter-subject variance is high (e.g., when studying high-level cognitive tasks such as recognition and navigation) or when only a limited amount of data per subject can be obtained (e.g., in one-shot experiments in recognition), large numbers of subjects must participate in the experiment. In practice, however, the number of subjects and stimuli in such studies is limited by practicalities like time and money constraints. Each subject has to be summoned, the instructions for the experiments have to be explained over and over again, a dedicated work-station has to be allocated for the subject (disabling the possibility of working with several subjects in parallel) and subjects have to be paid for their time. Such experiments, therefore, are usually performed on rather small numbers of participants (typically 10-50).

This paper suggests a way to overcome some of these limitations while benefiting from the new tool of VR simulation. Instead of bringing the subjects to the lab, we bring the experiment to the subjects. Making use of recent developments in web applications, in particular those that facilitate bi-directional interaction, we developed a web site on which a variety of spatial cognition experiments are provided. To simulate the 3D graphics used in many of our experiments on spatial perception and behavior, we crucially rely on VRML.

The advantages of using the web as a platform for our experiments are straightforward: the web enables us to reach a potentially huge number of subjects; these subjects use their own machines, instead of occupying limited resources at the lab; they don't require a dedicated researcher to explain the details of the experiment or to advise them throughout the process; and they don't have to be paid. The results of the experiments can be automatically accumulated, and even processed immediately, without intervention of the researcher. However, as we will discuss later, there are several disadvantages to such an approach, and the experiments have to be carefully designed so as to exploit the advantages of the new medium without undermining the scientific rigor of the results.

We are not the first to suggest using the net for conducting psychological and cognition studies. Hewson, Laurent and Vogel [5] describe methodologies for conducting psychological and sociological studies via the internet. Reips [9] discusses the use of the world-wide web for performing psychological experiments. A list of various experiments and studies in psychology that are available on the net can be found in the American Psychological Society webpage devoted to "Psychological Research on the Net" (<http://psych.hanover.edu/APS/exponnet.html>). However, most of this work is very preliminary and non of it uses VRML.

Various spatial cognition and navigation studies have been using virtual reality. May, Péruch and Savoyant [8] use a virtual environment to study the influence of topographical map alignment on navigation. Darken and Silbert [2] conducted an experiment to determine whether people use physical world wayfinding strategies in large virtual worlds. Distler [3] describes an implementation of a bicycle simulator to investigate orientation and navigation in a large-scale environments. Virtual reality is also used to study the ability of human subjects to search locations, to find shortcuts and novel paths, to estimate distances between remembered places, and

* galia.givaty@tuebingen.mpg.de

† hendrik-jan.veen@tuebingen.mpg.de

‡ chris.christou@tuebingen.mpg.de

§ heinrich.buelthoff@tuebingen.mpg.de

to draw sketch maps of the explored environment [4]. However, combining the power of the web with the use of virtual reality for spatial cognition and behavior experiments is a novel idea.

This paper concentrates on technical aspects of our approach. In section 2 we describe a particular experiment that was thus implemented. Details of the experimental and technical problems encountered and how they were overcome are given in section 3. In section 4 we give some preliminary results of the experiment. An evaluation of the application, listing its limitations and the class of experiments for which it is best used, is given in section 5. We conclude by suggesting ideas for further research.

2 DESCRIPTION OF THE EXPERIMENT

We describe in this section the application of our approach to a cognitive experiment that investigates visual encoding of objects. It is currently believed that humans form representations of the world by a hierarchical process of extracting and encoding information relating to objects shape and other attributes such as color, material properties, texture etc. [7]. Physiological evidence exists which suggests the human brain stores information regarding "where" an object is located separately from information about "what" particular properties an object has [10].

The purpose of this experiment is to obtain coarse indication of what mistakes people make when asked to reconstruct a particular conjunction of objects, their properties and their location in an environment. The kinds of mistakes made could give us further indication as to the form of human mental representation. The ability to isolate all but pertinent cues makes the use of Virtual Reality methods essential in this respect. The subject is first presented with a 3D display of a small room, in which several volumetric objects, each with different shape and visual properties, are randomly located (see figure 1). The subject is given ten minutes to stroll in this virtual world and memorize the arrangement of the objects. In the second stage of the experiment, a two-dimensional map of the room is presented, along with a set of cross-sectional representations of the objects, in all possible colors (see figure 2). The subject is required to reconstruct the arrangement of objects (according to their color, shape and position in the room).

Once the subject has placed all objects, the colors and shapes of the objects and their relative locations, along with the original (random) arrangement of the objects in the world, are recorded and stored for further processing. Additional background regarding visual perception studies can be found in [1].

Like many of our experiments, this one also necessitates a display of a virtual environment in which subjects are able to navigate. Usually, we are using dedicated hardware (Silicon Graphics Onyx2 IR) and high-level graphic software such as Multigen and SGI Performer, to model these environments and to perform the real-time rendering needed. However, once the experiments are mediated through the web, the virtual environments are designed using VRML.

The index page of the Tele-Experiment site lists the experiments currently available (<http://exp.kyb.tuebingen.mpg.de/web-experiment>). The experiment itself consists of a few web pages. The first page gives some motivation for the study and lists the software requirements for running the experiment (in particular, a VRML plug-in). The next page, which is presented only to users who visit the site for the first time, requests the subjects to fill a personal details form. Then, the subject gets detailed instructions for the experiment. Since the ability to recognize different colors and distinguish among them is essential for this experiment, the subject must first pass a pre-test that verifies this ability (see section 3.1). The next page contains two frames. The top one displays a VRML file, the virtual room, with a random arrangement of objects on its floor. The view type is set to WALK, which enables the subject to

move forward and backward in the room, turn left and right and change the viewpoint orientation vertically. The bottom frame displays the time left to examine the room.

The subject can stroll freely in this virtual world, and to indicate that all objects were found, s/he has to click on each of them. A click on the door of the room triggers the end of the memorizing part of the experiment and the display of the next page. This page uses a Java applet which displays a board of objects and the map of the room on which the subject has to place the objects as described above. After the experiment is over, the subject gets some feedback about his/her performances, as well as some additional information about the study.

3 IMPLEMENTATION

Several problems were encountered while adapting the experiment for the internet domain. We focus in this section on the problems and the methods used to overcome them.

3.1 Problems Inherent To Tele-Experimentation

Transferring the experiment from the controlled lab to the subjects' own environment results in three types of problems: retrieving personal details, relevant to the experiment, about the subject; controlling and monitoring the subjects performance; and motivating the subjects to conduct the experiment, keeping them concentrated throughout the process.

There are many differences between running a fully controlled experiment in the lab, and having subjects perform the experiment in their own uncontrolled environment. In the former case, all the information about the conditions in which the experiment is performed, and personal information pertaining to the subjects is known. Undertaking an experiment over the web enables the subject to stay absolutely anonymous. Relevant details such as the age, sex or education of the subjects is unavailable. Similarly, important parameters such as the hardware on which the experiment is performed, the operating system used, the web browser or the VRML plug-in, which are far more diverse than in the case of a lab experiment, cannot always be retrieved.

Obviously, these problems cannot be solved completely. We try to retrieve as many parameters as possible: First, we ask each subject to fill a personal details form, containing data such as age and gender, before running the experiment. Second, we record the user's IP address, operating system and browser information. Clearly, this might not be sufficient; however, the distortion of the results caused by the missing data is usually compensated for by the large number of participants.

Some experiments can have additional requirements from the subject. For example, the experiment described in section 2 requires a normal color vision. Color-blind subjects, as well as users without color displays, must be eliminated from participating in the experiment. To ensure this, a short pre-test is used. The user is shown 5 buttons with different colors, ordered randomly. To start the experiment, s/he should press the red one. Pressing any other button leads to a page with an error message, and prevents the user from running the experiment.

Another problem is that there can be no control over the behavior of subjects during the experiment: they can perform it partially, interruptedly or multiple times. Such parameters might affect the validity of the results. To minimize the chances of a user participating more than once in a single experiment, we use cookies stored on the user's own system. Since checking cookies is possible only from within JavaScript, participating in the experiments is allowed only with browsers that support JavaScript. The time given to the subject to perform the experiment is limited, and if the task is not

completed within the given period, the results are discarded. Furthermore, we record the time of “important” events (like clicking on an object) during the experiments, using VRML’s TimeSensor node. This information helps the researcher to monitor the subject’s behavior.

The last problem is to motivate subjects to take part in the experiment. Subjects in the lab are paid for their time, and thus can be requested to perform tedious, sometimes boring tasks, for long durations. On the web, however, the experiment must be kept short and interesting or it will not appeal to potential subjects. Furthermore, a boring experiment might cause potential subjects to abort their participation and never come back for other experiments. However, navigating through a virtual scene is an attractive experience which itself motivates subjects to participate in the experiment.

3.2 Browser/Plug-in Parameters Control

When navigating in a VRML world, the user has control over various global parameters. The first and most critical one for this application is the navigation paradigm. Experiments such as the one described here, which focus on human navigation capabilities, must simulate a “real world” environment, and in particular, limit the subjects to WALK view. When using other views, e.g., the EXAMINE view, the subject can obtain additional information about the scene which renders the experimental results invalid. Browsers support ANY, WALK, EXAMINE, FLY, and NONE navigation types, and an ability to restrict the view paradigm is part of the VRML 2.0 specification. However, in some browsers (e.g., WorldView v2.0) it is impossible to disable a subset of the possible views. Furthermore, it is impossible to retrieve the current view type, so the researcher cannot know whether the subject has used a different view type than the one designated for the experiment. This problem, too, cannot be fully solved. As a partial solution, the user’s browser name and version are recorded. If the performance of subjects using browsers that do not support view types restrictions, are found to be significantly different from other subjects, their results can be ignored.

A similar problem is caused by the ability to turn off the collision detection mechanism. Again, this can give the user view points which are initially blocked. Another parameter that influences the appearance of the world is the field of view. The horizontal field of view can be set by the VRML author, but there is no control over the vertical one. Changing the size of the window that displays the VRML world results in distortions and different vertical fields of view.

There are many other, less important, factors, over which users might have control. For example, a user can adjust the navigation speed, control the level of details, texture display, transparency, headlights, etc. The actual values of most of these parameters are unavailable to the experiment designer.

A somewhat radical solution for the above problems is writing our own VRML 2.0 plug-in. Such a solution will enable us to have full control over the above parameters. Unless we decide to do so, we must settle for the fact that different users might have different visual data for the experiments.

3.3 Script Nodes Support

There is a variety of VRML 2.0 plug-ins for various operating systems and browsers. One of the most important differences among different plug-ins is their support of Script nodes. The most commonly supported languages are: VRML Script, Java Script and Java. However, some plug-ins do not support all of them. For example, SGI Cosmo Player v1.0.2, a commonly used plug-in for Netscape on IRIX systems, supports only VRML Script. To enable as many participants as possible to take part in the experiments, it was important to have as many plug-ins as possible supporting the

application. Since most browsers support at least one of the two script languages (VRML Script or Java Script), we restricted the implementation to these two languages only. Each Script node has two URLs on its *url* field; one for each language. Although Java is far more powerful than both of these script languages, we decided not to use it.

3.4 Transferring Data To The Host

Apart from obtaining the final results, an essential part of most spatial cognition experiments is retrieving information regarding the subject’s behavior *during* the experiment. The moves of a subject in a virtual world, the reaction times, the progress speed etc., are relevant to many experiments. The experiment described above, for example, involves a random collection of objects in the world. Since the particular arrangement is only known at run-time, this information must be transferred to the host along with the results. Information transfer from the client (the subject’s machine) to the server (the experiment site) is not part of the standard VRML language. Such features must be implemented using Script nodes. While it would have been possible to achieve such functionality using *URLConnection* class or *sockets* in Java, this solution was ruled out due to reasons discussed above (section 3.3).

As an alternative, we collect the necessary information in a String object using a script language (VRML Script or Java Script). This necessitates a concise encoding of the information, but is nevertheless possible. The string records the initial settings, and additional information (e.g., the times of mouse-click events) is appended to it as the subject explores the VRML-world. Since the problem of transferring information to the host only exists with VRML (it is always possible to send the information to the server from within an HTML document), we invoke an HTML page from the VRML environment, using the *loadURL()* function of the Browser object. The contents of the string are appended to the URL name, in a similar fashion as the use of *CGI* with the *GET* method. The invoked HTML page retrieves the value of the string from its own URL, using Java Script, and transfers it to a Java applet that is executed in this page. The applet adds the information to other data that are collected in this HTML page, and sends all the information to the server.

Multiple-page experiments raise another problem. Data are produced in various pages and must be combined together to create a complete picture. Combining the information can be done either during the experiment execution or after the data are sent to the host.

In the former case, the information is transferred from one page to the next, using strings, as described above. We use this method to transfer information from the VRML page to the next one. In the latter case the data are sent to the server in segments from different pages. In order to be able to fit the different parts together, a unique ID, corresponding to the subject performing the experiment, is attached to each of the different segments. When all the data are received by the server, they are combined to one record, associated with a particular subject. This method is used, in the described experiment, to connect the personal data of the subjects that are available in the introductory page, with the results of the experiment that are only known later.

4 PRELIMINARY RESULTS

The experiment is still in its infancy but is now being advertised in an increasing number of places on the internet. To date only 34 people have participated. The average time these participants spent observing the room was 252 seconds (4.2 minutes) with a minimum of 39 seconds (rather fast) and a maximum of 600 seconds (the maximum allowed).

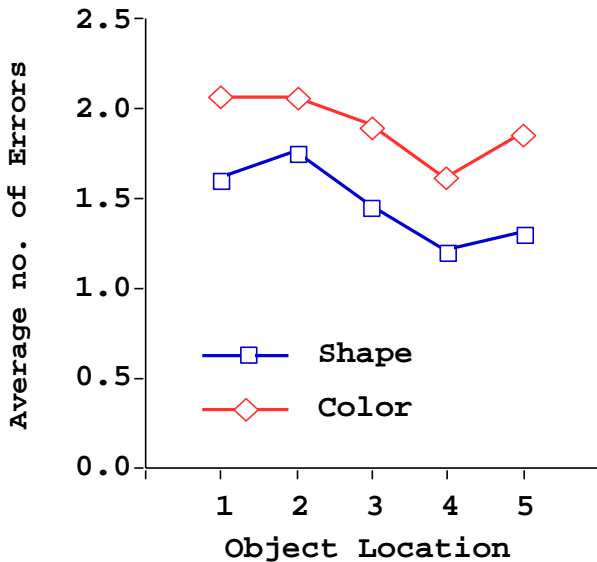


Figure 3: Average number of errors as a function of the objects location within the VRML environment

The task of each observer was to reconstruct the spatial organization of the VRML environment, by putting the correct shape and the correct color on every position. We compare the number of mistakes participants made in determining the color of an object with the number of mistakes made in determining its shape. The graph in figure 3 displays the average number of errors, over all subjects, made in color and in shape as a function of the object location (marked with numbers 1..5). It is still early-days to draw firm inferences from the data collected, but nevertheless, the graph shows that consistently fewer errors were made for shape than for color. On average, over all 5 spatial locations and over all subjects, participants made 1.5 errors for shape placement and 1.9 errors for matching color (out of 5 decisions). Complete guessing would, on average, result in 4 errors for both shape and color. The difference between the shape and color dimensions is significant only with a probability of error of 0.13. That is, we have a 13% chance of making the wrong conclusion. The problem is that the differences between the two dimensions are quite small and the inter-subject variability is quite high (our initial expectation). We clearly require many more participants to make any firm conclusions either way, which is exactly why we do this experiment on the web.

5 EVALUATION

Tele-Experimenting involves a subject at a remote location using his or her own hardware and software to run an experiment. This situation differs from a classical experiment performed in the laboratory in at least two major ways: (1) the researchers have no direct contact with the subject; and (2) they cannot control the hardware and software that the subject is using for the experiment. Whether or not Tele-Experimenting will be accepted as a valid and useful research method by the scientific community decisively depends on how we deal with those two aspects.

A few general guidelines for doing Tele-Experiments can be formulated. First of all, careful selection of experiments can avoid many problems. Those experiments that critically rely on parameters that cannot be precisely controlled in Tele-Experimenting (such as stimulus contrast, graphics quality, etc.) should obviously be

avoided. Secondly, one should try to retrieve as many parameters about the subject and the hardware and software s/he is using as possible. The idea is that if you cannot control a parameter then at least you should try to monitor its value. Knowing the value of parameters enables one to make selections afterwards. Thirdly, one should always consider to run the same experiment under controlled laboratory conditions with a small number of subjects to provide comparison data. When more and more of such comparisons are being collected we will eventually be able to address some issues pertaining to the question of the general validity of Tele-Experiments.

The class of experiments that can be conducted as Tele-Experiments is restricted by a few factors. One important factor is time: the experiment should be fast enough to keep the anonymous subject interested and concentrated. Here it makes more sense to borrow ideas from the game industry than to use the standard psychophysical procedures. Typically this means that extensively repetitive experiments or multi-stage experiments are not suited for Tele-Experimenting. A second factor is subject anonymity. Although we can ask subjects to reveal some information about themselves, e.g., gender and age, we have no guarantee whatsoever that this information is correct. Moreover, we cannot control or monitor whether they are doing the experiment alone or with the whole family. Experiments that are too sensitive to this factor will produce invalid results. A third important factor relates to the stimulus specification. Parameters like contrast or brightness of the imagery cannot be precisely controlled or even just monitored. Experiments that do depend on such parameters, and many classical psychophysical experiments on early-vision do, should not be conducted as Tele-Experiments. The last important factor that should be kept in mind is the type of task the subjects have to perform. Fast, attractive, non-repetitive, intuitive tasks are preferred. The experiment described in this paper features a map-construction task, rather than a large set of questions about the layout of the scene, precisely for this reason. In summary, Tele-Experimenting requires careful planning of the experimental design, procedure and tasks that researchers want to use.

6 CONCLUSION

We have described in this paper an application of VRML-based multimedia to the domain of cognitive experiments. The advantages of our approach are straightforward: it enables us to access a potentially huge number of participants in our experiments, thus overcoming the bottleneck of such experiments in the past. We listed several of the limitations of our approach, suggesting some restrictions on the experiments to which it is applicable. A detailed presentation of a particular experiment, crucially using VRML to simulate a virtual environment in which subjects are required to navigate, was presented.

While the paper describes many solutions to both experimental and technical problems of using VRML in this domain, several other problems remain unsolved and are the subject for further investigations. In particular, we would like to be able to have a continuous data stream out of the VRML world towards the host. This would allow us to get more data about subjects' behavior in the virtual environment, e.g., about their exploration behavior. Another important issue is the use of small tests to monitor properties of the hardware and the subject. The experiment described in this paper had two of these tests: the color test in the beginning is one of them. The other one is built around subjects having to click on all of the items in the VRML-world: this ensures they could see and distinguish these objects and it also guarantees that they were able to use the browser to move through the environment. Similar tests can be designed to test for the availability of audio output, for the rendering quality of the visual scenes, etc.

7 ACKNOWLEDGMENTS

We would like to thank Stephan Braun and Philipp Georg, who helped us with some of the technical issues described in this paper.

References

- [1] V. Bruce and P. R. Green. *Visual Perception: Physiology, Psychology and Ecology*. Lawrence Erlbaum Associates, 1990.
- [2] Rudolph P. Darken and John L. Silbert. Wayfinding Strategies and Behaviors in Large Virtual Worlds. *Proceedings of ACM SIGCHI*, pages 142–149, 1996.
- [3] Hartwig Distler. Psychophysical Experiments in Virtual Environments. *Virtual Reality World '96 Conference Documentation*, 1996.
- [4] S. Gillner and H. A. Mallot. Navigation and Acquisition of Spatial Knowledge in a Virtual Maze. Technical Report 45, Max-Planck-Institute for Biological Cybernetics, February 1997.
- [5] C. M. Hewson, D. Laurent, and C. M. Vogel. Proper Methodologies for Psychological and Sociological Studies Conducted via the Internet. *Behavioral Research Methods, Instruments and Computers*, 28:186–191, 1996.
- [6] Roger E. Kirk. *Experimental Design: Procedures for the Behavioral Sciences*. Brooks/Cole Publishing Company, 1995.
- [7] D. Marr and H.K. Nishihara. Representation and Recognition of the Spatial Organization of Three Dimensional Shapes. *Proceedings of the Royal Society of London, Series B*(200):269–294, 1978.
- [8] Mark May, Patrick Péruch, and Alain Savoyant. Navigating in a Virtual Environment with Map-Acquired Knowledge: Encoding and Alignment Effects. *Ecological Psychology*, 7(1):21–36, 1995.
- [9] Ulf-Dietrich Reips. *Psychological Experimenting on the Internet*. In B. Batinic (Hrsg.): *Internet Für Psychologen*, chapter 11, pages 245–265. Hogrefe: Göttingen, Germany, 1997.
- [10] Leslie G. Ungerleider and James V. Haxby. ‘What’ and ‘Where’ in the Human Brain. *Current Opinion in Neurobiology*, 4:157–165, 1994.