

Doron Friedman*

Department of Computer Science
University College London
London, UK

Robert Leeb

Laboratory of Brain-Computer
Interface
Graz University of Technology
Graz, Austria

Christoph Guger

g.tec—Guger Technologies OEG
Graz, Austria

Anthony Steed

Department of Computer Science
University College London
London, UK

Gert Pfurtscheller

Laboratory of Brain-Computer
Interface
Graz University of Technology
Graz, Austria

Mel Slater

Department of Computer Science
University College London
London, UK
and
ICREA Universitat Politècnica de
Catalunya
Spain

Navigating Virtual Reality by Thought: What Is It Like?

Abstract

We have set up a brain-computer interface (BCI) to be used as an input device to a highly immersive virtual reality CAVE-like system. We have carried out two navigation experiments: three subjects were required to rotate in a virtual bar room by imagining left or right hand movement, and to walk along a single axis in a virtual street by imagining foot or hand movement. In this paper we focus on the subjective experience of navigating virtual reality “by thought,” and on the interrelations between BCI and presence.

I Introduction

Virtual reality (VR) research is continuously striving towards natural and seamless human-computer interfaces, and the existing interfaces for locomotion through virtual environments (VE) are still not satisfactory. Typically, participants navigate by using a handheld device such as a joystick or a wand. They are then exposed to conflicting stimuli: the world around them seems as if they were moving, but they feel that their body is stationary. This results in a reduced sense of being present in the VE (Slater, Usoh, & Steed, 1995), and is one of the causes of simulation sickness (Hettinger & Riccio, 1992). Slater, Usoh, and Steed investigated a method that allows participants to walk in VR by walking in place; people using this method reported a higher sense of presence on the average than those who locomoted using a pointing device. In a later experiment (Usoh et al., 1999) walking in place was compared with really walking, and in terms of the reported sense of presence, the results were not much different. Rather than walking in place, what would it be like if we were able to navigate a VE by merely imagining ourselves walking?

In the vision portrayed by science fiction, the brain is the ultimate interface to VR. While we are still far from an integration of the brain and VR, BCI makes it possible to start exploring this possibility. Electroencephalogram (EEG)-based BCI research is aimed at helping individuals with severe motor deficits to become more independent (Wolpaw, Birbaumer, McFarland, Pfurtscheller, & Vaughan, 2002). It has been shown (Pfurtscheller & Neuper, 2001) that it is possible to identify a few mental processes using electrodes attached to the scalp, that is, the imagination of various predefined movements, from online EEG signals. Such thought-related EEG changes have been trans-

formed into a control signal and associated with simple computer commands (i.e., cursor movement). The Graz-BCI paradigm (Pfurtscheller et al., 2003) is based on motor imagery, which makes it a natural candidate for thought-based navigation in VR.

We have set up a system that connects the Graz-BCI to a highly immersive four-sided CAVE-like (Cruz-Neira, Sandin, DeFanti, Kenyon, & Hart, 1992) system. Three subjects, trained at Graz, Austria, used the BCI in two navigation tasks: rotating in a virtual bar room and moving in one dimension along a virtual street. Each subject had two sessions in both VR conditions. Following the VR sessions they were asked to fill in questionnaires, and non-structured interviews were conducted.

The experiment showed that BCI could be used to control locomotion events in a CAVE-like setting. We were also able to compare our experiments in the CAVE with similar experiments carried out earlier with a head-mounted display (HMD) (Leeb, Scherer, Lee, Bischof, & Pfurtscheller, 2004). The BCI techniques and BCI accuracy results will be detailed in a separate paper, and only explained here briefly. In this paper we are more interested in the participant's experience: given that navigation by thought is possible, what is it like? How does BCI interact with presence? What are the constraints and how can they be overcome? What design decisions need to be made, and how do they affect the experience? Given that our results indicate that VR can improve BCI accuracy, these questions have practical implications.

2 Background

Previous research has established that a BCI may be used to control events within a VE, and some research has also been done in immersive systems. Nelson, Hettlinger, Cunningham, and Roe (1997) were interested in BCI as a potential means for increasing the effectiveness of future tactical airborne crew stations. They have investigated the usage of CyberLink, an interface that uses a combination of EEG and electromyography (EMG) biopotentials as control inputs, in a single-axis continuous control task. The participants

used the CyberLink interface to navigate along a predetermined flight course that was projected onto a 40-foot diameter dome display. Continuous feedback was provided by a graphical heads-up display (HUD). Participants were not given any BCI instructions or training. Scores of effective task performance gradually increased with training and reached an average of 80% success.

Middendorf, McMillan, Calhoun, and Jones (2000) harnessed the steady-state visual-evoked response (SSVER), a periodic response elicited by the repetitive presentation of a visual stimulus, as a communication medium for the BCI. SSVER can be used for BCI in several ways. In this experiment two methods were employed; one of them was tested with a flight simulator. In this method operators were trained to exert voluntary control over the strength of their SSVER. One of the conditions involved controlling a flight simulator, where the roll position of the flight simulator was controlled with BCI. The simulator rolled right if 75% or more of the SSVER amplitude samples over a half second period were higher than some threshold, and left if most of the samples were lower than another threshold. Most operators were able to reach 85–90% of success after 30 min of training.

Bayliss and Ballard (2000) used the P3 evoked potential (EP), a positive waveform occurring approximately 300–450 ms after an infrequent task-relevant stimulus. They used HMD-based VR system. Subjects were instructed to drive within a virtual town and stop at red lights while ignoring both green and yellow lights. The red lights were made to be rare enough to make the P3 EP usable. The subjects were driving a modified go-cart. Whenever a red light was displayed, data was recorded continuously from –100 to 1000 ms. Results show that a P3 EP indeed occurs at red and not yellow lights, with recognition rates that make it a candidate BCI communication medium.

In further research Bayliss (2003) continued exploring the usage of the P3 EP in VR. Subjects were asked to control several objects or commands in a virtual apartment: a lamp, a stereo system, a television set, a Hi command, and a Bye command in several nonimmersive conditions and with an HMD. Using BCI, subjects could switch the objects on and off or cause the ani-

mated character to appear or disappear. The BCI worked as follows: approximately once per second a semi-transparent sphere would appear on a randomly selected object, for 250 ms. Subjects were asked to count the flashes on a specific object (to make the stimulus task-related, as P3 requires). An epoch size from 100 ms (before the stimulus) to 1500 ms was specified. Text instructions in the bottom of the screen indicated the goal object. The subject had to count the flashes for that object only. The subject was given a visual feedback when a goal was achieved, that is, when a P3 event was recorded when the target object was flashing. Subjects were able to achieve approximately 3 goals per minute. Bayliss found no significant difference in BCI performance between VR and a computer monitor. Most subjects preferred the VR environment; all of them liked the fixed-display condition (looking through a fixed HMD) the least.

All the research described so far is based on several types of visually evoked responses. These methods typically force a visual task on the subject, which might be unnatural. Our research is based on a different BCI paradigm: the Graz-BCI paradigm is based on motor imagery, in a VR rotation task. Leeb et al. (2004) have used a fixed-display HMD setting, which means the subjects actually experienced a limited form of VR. They report BCI performance success rates of 77–100%. This was the first step in the research reported here.

3 The Experiments

3.1 The System

In order to carry out the navigation by thought experiments we had to integrate two complex software and hardware systems: the BCI and the CAVE-like VR system. A system diagram appears in Figure 1 and is explained below.

The experiments were carried out in a four-sided ReaCTor (CAVE-hybrid; Cruz-Niera et al., 1992) system, which is driven by an Onyx IR2 with four graphics pipes. Participants were head tracked using a wireless Intersense tracker. The applications were implemented on top of the DIVE software (Frecon, Smith, Steed,

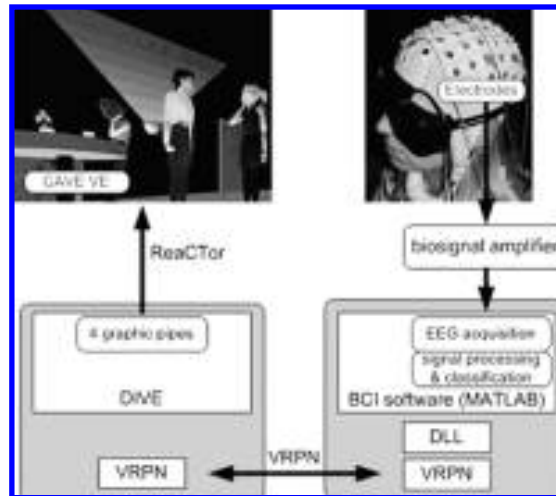


Figure 1. A diagram of the BCI-CAVE integrated system.

Stenius, & Stahl, 2001; Steed, Mortensen, & Frecon, 2001).

The Graz-BCI system consists of a bio-signal amplifier (g.tec, Graz, Austria), a data acquisition card (National Instruments, Austin, TX, USA) and a standard PC running Windows 2000. The signal processing is based on Matlab and Simulink (Mathworks, Natick, MA, USA).

The communication between the PC running the BCI and the VR host is done using a communication system called virtual reality peripheral network (VRPN). VRPN provides synchronization and logging of multiple data channels, and has built-in support for many VR devices.

On the PC it communicated with the BCI Matlab-based software via a dynamic-link library (DLL). On the UNIX machine controlling the VR, a VRPN plug-in for DIVE was implemented. VRPN uses UDP to establish connection and TCP for sending messages over the network. The main challenge was to balance the CPU time between the communication and the rest of the system: the BCI computation on the PC, and DIVE VR rendering on the UNIX machine. We were interested in receiving 20 updates per second with a minimum delay time, so that the feedback for the BCI decisions would be as smooth as possible.

3.2 Brain-Computer Interface

The experiment included three subjects: one female and two males. All subjects were familiar with the Graz-BCI (Pfurtscheller et al., 2003) over a period of between four months and two years. In addition, they were specifically trained for this experiment by performing identical tasks in similar VEs with an HMD. The VEs were not identical: in the training phase the subjects had to rotate an auditorium, whereas in the CAVE experiment the subjects had to rotate a virtual bar. In the walking task we used two different models of streets. Two different experiments were performed. One required rotating to the left or to the right, inside a virtual bar, by imaging a right or left hand movement, and the second experiment required moving forward along a virtual street by imaging a foot movement.

In all experiments the subject was sitting on a comfortable chair in the middle of the CAVE (see Figure 2). Three bipolar EEG channels (electrodes located 2.5 cm anterior and posterior to C3, Cz, and C4, respectively, according to the international 10–20 Hz system) were recorded with a sampling frequency of 250 Hz. The logarithmic band power was calculated sample-by-sample for 1-s epochs in the alpha (10–12 Hz) and beta (16–24 Hz) bands of the ongoing EEG and classified by a linear discriminant analysis (LDA). The LDA classification result was used as a binary control signal and sent via the VRPN to the CAVE system to modify the VE. The LDA classifier used in these experiments was calculated offline from data recorded previously using similar VEs in an HMD experiment (Leeb et al., 2004).

Each subject participated in two sessions on two consecutive days and each session included four feedback runs. Each run consisted of 40 trials (20 left and 20 right cues, in the case of the rotation experiment and 20 foot and 20 right-hand cues for the walking experiment) and lasted about seven minutes. The sequence of right/left or foot/right cues was randomized through each run. Depending on the affiliation of the acoustic cue, the subject was instructed to imagine a left hand or right hand movement in the bar-rotation experiment; or to imagine right hand movement or the movement of either foot in the street-walking experiment.



Figure 2. Two images of a subject connected to the BCI, wearing shutter glasses, in the virtual bar VE in the VR CAVE room. Using only his thoughts, the subject can rotate the VE to pay attention to different characters in the bar.

3.3 Experimental Setup

BCI can be realized in an externally-paced mode (synchronous BCI) or in an internally-paced mode (asynchronous BCI). In synchronous BCI, specific mental patterns have to be generated in response to an external event, that is, changes in brain activity are tracked over a predefined time window. In asynchronous BCI the EEG has to be analyzed continuously. We have used a synchronous BCI which is more limited, but more

reliable. This has a disadvantage compared with traditional VR navigation devices.

Classification of the signal can start immediately after the trigger, but the optimal classification point varies between individuals, and is typically at least two seconds after the trigger (Krausz, Scherer, Korisek, & Pfurtscheller, 2003). This delay is too long for a user interface (UI) feedback, thus we prefer to provide visual feedback immediately after the BCI decision. The system sent the classification result every 52 ms approximately, over a period of 4.16 s after the trigger. Continuous feedback was provided in the form of continuous viewpoint translation or rotation.

The course of events is as follows: the application decides that a navigation decision is required, and DIVE sends a request over the VRPN network. This request, along with all other network events, is logged with an accurate timestamp for post hoc analysis. The request is intercepted by the VRPN component running on the BCI PC. It communicates with the DLL, which then sends a request to the Matlab BCI software. The BCI software makes a random decision about the navigation decision required by the participant, and initiates an auditory trigger accordingly. The BCI software immediately starts analyzing the EEG signals, and makes a classification decision approximately every 52 ms. A binary value is passed back to the DLL, and then over the VRPN network. The DIVE VRPN plug-in, on the UNIX host, intercepts this event and feeds it into the application. The application then changes the VE to reflect the participants' rotation or translation.

We selected two VEs for simple navigation tasks. The first task was rotation: the VE depicts a virtual bar (see Figure 2). It is populated by four virtual characters and a virtual barman. Originally, the virtual characters talked to the participants when their head gaze and the participant's gaze crossed (this is possible since the participant is head tracked), and the bar included background music and chatter. Eventually we decided to remove the background music and the speech audio, because the BCI triggers were auditory, and we wanted the subjects to hear them clearly. We left the background chatter, and that was the first time these BCI subjects were ex-



Figure 3. A BCI subject in the virtual street scene.

posed to an experiment with an audio track, other than the trigger signals.

The size of the virtual bar is not much larger than the size of the physical projection room, and there is no need for the participant to move around the VE; it is enough for them to look around. This environment is, thus, suitable for a simple first experiment, in which participants operate in the VE by rotation only.

The BCI sent 80 classification results over a period of 4.16 in fixed intervals. These results were used by the VR to provide continuous feedback; the total rotation per trigger could be up to 45°, and for each BCI update the VE rotated by 0.56°.

The second task was to simulate walking: the VE depicts a long main street with various shops on both sides (see Figure 3). Some of the shops could theoretically be visited inside, but in the BCI experiments we carried so far the participants could only go straight, so they could not enter the shops. The street is populated with some animated characters that walk along the street. The characters are programmed to avoid collisions.

BCI control was as follows: if the trigger indicated that the participant could walk, the participant had to imagine foot movement to move forward. If the trigger indicated that the participant should stop, they had to imagine right hand movement. If the classification indicated hand imagery when ordered to walk, the partici-

pant would stay in place. If the classification indicated foot imagery when ordered to stop, the participant would go backwards. This “punishment” was introduced so that the subjects will not be able to adopt a strategy of always imagining their foot. The same timing as in the rotation experiment was used, but instead of rotating, the BCI updates moved the subject by 0.05 distance units, so the total distance per trigger could be up to four units.

4 Results

In this paper we are mostly interested in evaluating the participants’ subjective experience. We will briefly report the overall BCI performance results; details will be reported in a separate paper. The three subjects (L1, O3, S1) who participated in this study were able to rotate the bar with a relatively high level of success: the average performed rotation of one run to the left was -21° and to the right 22° ; resulting in a BCI performance from 80 to 100%. Two of the three were also able to navigate the street. One of the three subjects (S1) was not able to stop properly using hand motor imagery, and was thus replaced with subject J8.

The cumulative distance traveled by the subjects in the street scenario can be used as a performance measurement of the experiment. The maximum achievable mileage would be 80 distance units and the result of a random session would be 0 units. Note that this cumulative performance measure is different than the BCI performance. Recall that for each subject there is an optimal duration for classification within the BCI epoch: BCI performance can be assessed taking into account only the classification achieved in this optimal point in time. The mileage performance takes into account classification throughout the whole 4.16 s epoch.

Figure 4 compares the performance for each session of the three different conditions for the three subjects. The results performed in the CAVE are displayed in the right columns, the performance achieved with the HMD is displayed in the middle columns, and the performance of the standard BCI without using the VR as a feedback medium are displayed in the left columns.

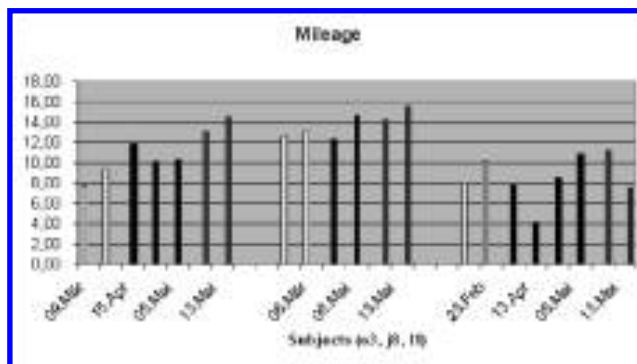


Figure 4. BCI performance as measured by virtual mileage. The details for the subjects O3, J8, and I1 appear from left to right, respectively. Within each subjects, three conditions are displayed from left to right, respectively: traditional BCI, HMD, and CAVE-like VR.

The mileage for the standard BCI (left columns), were simulated offline to be able to compare the results. In all cases the performance in the HMD condition was significantly better than the performance in the monitor condition and in all cases the performance in the highly immersive condition was better than all other conditions. This was statistically significant in two out of three subjects: subject O3: $F_{2,21} = 22.08$; $p < .00001$ and subject J8: $F_{2,18} = 9.63$; $p < .002$.

After completing a session, the subjects were asked to fill in the Slater-Usch-Steed presence questionnaire (Slater, 1999), and then a non-structured interview was conducted. This was carried out with three subjects after the bar VE and with two subjects after the street VE. Based on questionnaire and interview data, we can evaluate several aspects of presence: overall sense of presence, body representation, and social presence. In addition, we can look at other interactions between BCI and VR.

Subjects were BCI trained for long durations (of four months up to two years), and they were specifically trained for the same task. The subjects were very limited in their actions: they are trained not to move or blink during the BCI feedback epochs. Since our experiment included 4s of BCI feedback followed by approximately 5s rest until the next trigger, they were asked not to move throughout the whole experience. These constraints posed by the BCI make it difficult to evaluate

the sense of presence, as well as the overall experience. It is, however, possible to compare the experience of traditional BCI with the experience of BCI in a highly immersive VR.

4.1 Overall Presence

Appendix A provides the details for five different questions from the questionnaire that relate to overall presence, and that were applicable to the BCI experiments. The number of subjects is too small to allow statistical analysis, but the results are consistent: there is a difference among subjects: L1 reported high presence, S1 reported medium presence, and O3 reported low presence. There seems to be no significant difference in levels of presence between tasks.

In the post-questionnaire interviews subjects reported that concentration on the BCI task interacted with presence. Since gradually the BCI control became more automatic, they could gradually be more absorbed by the VR and feel more present. This description is different from what most subjects who experience the same VEs without BCI report: initially they feel a high sense of presence, but this gradually drops as they realize the limitations of the VR (Garau et al., 2004).

All subjects report that the CAVE is more comfortable than the HMD. One subject reported that the wide field-of-view made him feel as if the landmarks were all around him, more like in the real world than in a typical BCI session; this may imply that spatial presence may facilitate BCI. The subject mentioned he did not experience this with the HMD setting.

Subjects report a conflict between what they wanted to do during the experience and the BCI task for which they were trained. One subject wanted to reach his hands for the beer or talk to the characters in the room; he said: "It was like a little voice in my head saying 'try this, try this,' but I know I am not allowed to."

Subjects noted that the bar room had two areas: the virtual characters concentrated in one area, whereas the other side of the room was empty, and didn't even contain furniture (only a mirror ball). One subject said BCI control was easier for her in that area, because it was less

distracting. For another subject BCI control seemed more difficult in the empty space, because there was no clear spatial information.

The audio chatter was difficult for one subject. It was distracting, but also annoying because of repetition. He noted that the visuals repeated too, but this was not annoying or distracting. The difference could be due to the fact that the audio was a repeated loop (its duration was one minute), whereas the character's gestures are pseudorandom. Also, the subjects reported that part of the audio track included a very distinct laughter, which stood out when repeating itself.

4.2 Body Representation

Previous work has suggested the very important role of perception of the body within virtual reality (Slater, Usoh, & Steed, 1994): the more that the body is used naturally, and the more that it is anchored into the VE, the greater the chance for presence. Such research refers to the real body, which is perhaps represented by a virtual avatar within the VR. Our case is different; BCI may be considered a (very unusual) extension of the body; it is thus interesting to learn about the subjects' sensation of their own body in the experience. Subjects reported that the experience felt natural compared to other BCI experiments. One subject noted that she felt as if her whole body were rotating. Another one asked whether he felt as if his body was actually rotating, answered: "No, it was more like in a dream—you move but you do not feel your body physically move. And just like in a dream—at that moment it seems real."

The interaction seemed more natural than traditional BCI to all subjects, even though the mapping between the motor imagery and the application functionality was not perfect: body rotation was mapped into hand movements, walking was mapped into arbitrary feet movement, and stopping was mapped to right hand movement. Yet all subjects reported that rotating and moving based on hand or feet imagery seemed quite natural, and that this was very different from the typical BCI training setting.

4.3 Social Presence

The two VEs included animated characters. It was thus possible to evaluate the subject's sense of social presence, or to what extent the subjects felt as if they were in a socially populated VE. Appendix B details a subset of the questionnaire questions related to social presence; however, unlike the overall presence questions, the results do not seem to be consistent even within subjects.

Subjects did not generally report a high level of social presence, though they certainly paid attention to the virtual characters. One subject commented that the characters in the bar were used as landmarks for the rotation, which made them be treated as inanimate objects. Another one, when asked if he felt the characters were real, replied: "Yes, but not exactly. It was as if I am a space explorer who just met some aliens. They look humanoid, but they behave different; as if they were some other life forms." Further research is needed to explore the effect of using an interface such as BCI on social presence.

5 Discussion and Future Work

The highly immersive (CAVE) condition achieved the highest level of BCI accuracy, compared to less immersive conditions. Further research is needed in order to understand why this happens: is this because participants are more motivated in highly immersive VR, or are there additional factors? Does presence play a role here?

Moreover, all subjects liked the CAVE setting more than the HMD, and both were very much preferred over BCI training on a monitor. Novelty could also be an important effect. The main reason given by the subjects for preferring the VR was that they provided motivation. Specifically, the street VE was treated as a sort of racecourse: subjects wanted to get further away in the street, and further than other subjects in previous sessions. An interesting comment was made by one of the subjects: motivation seems to greatly improve BCI per-

formance, but too much excitement might have a negative impact, as it makes it harder to concentrate on the BCI control. It may be interesting to explore this in follow-up research, and to isolate the relative contributions of presence and motivation in BCI performance.

Previous research has demonstrated that BCI-based control of VR is possible, and we were able to repeat this result in the CAVE-like setting. In addition to the feasibility of BCI control in the CAVE, the main lesson learned is that typical BCI procedures are, at the moment, too prohibitive to become a natural UI. Given that VR seems to be a promising counterpart to BCI, we may ask: what needs to be done in order to make BCI control of VR a positive experience for a wide range of participants?

Thus far we have not been able to demonstrate a realistic scenario of navigating a VR by thought; only subsets of the navigation task were demonstrated. Even in these simplified tasks, we had to remove substantial parts of the VE functionality to make it easier for subjects to concentrate. Both VEs did include significant visual input, and the bar included a background audio track as well. Future research will need to go beyond adapting the task to the constraints posed by the BCI, and assess BCI performance in a wide range of rich VEs with various types of tasks and interactions.

Even in the scope of these limited tasks, there are several limitations that come up. First, we note that for the classification to reach a high level of accuracy, the subjects need to be trained over many sessions, typically over a few days or weeks. Even then, the accuracy is seldom 100% (Pfurtscheller & Neuper, 2001), as is expected from traditional devices. In our case, note that subjects were not only highly trained to use BCI in general, but were also extensively trained for the specific task!

BCI poses physical constraints on the subject that might make it unacceptable for some applications. Subjects need to sit still, and during the BCI control epochs are trained to stop blinking, or move their muscles.

Another problem with BCI is that the reliable methods are typically trigger-based (synchronous).

This has a great disadvantage compared with traditional navigation devices. The best we can do in synchronous paradigms is to incorporate the BCI cues into the environment, in a manner that will be as seamless as possible.

Further difficulties are specific to the motor imagery paradigm, and may be avoided using other BCI techniques. Using the motor imagery paradigm, the classification is optimal only a few seconds after the cue. As previously explained, such delay is not acceptable for a UI, so we introduce continuous feedback. This has a penalty: the VE functionality does not depend on the optimal BCI decision.

Our aim is to continue this research towards free-style BCI-controlled navigation, in which the participants will be free to explore the VE rather than required to follow cues. Interestingly, such freestyle control of BCI in VR was never attempted, since researchers typically design experiments aimed only at measuring BCI accuracy. We might still be able to evaluate BCI performance by looking at overall task performance, where the task depends on navigating the VE. In addition, we want to try alleviating some of the constraints posed by the BCI, and try to make BCI a more natural interface for navigation.

The mapping between the recognized thought-related EEG patterns and the VR functionality needs to be further explored. We hope to learn whether a more natural mapping improves BCI accuracy or learning rates, and to what extent can participants adapt to less intuitive or even counterintuitive mappings. This is important because the number of thought patterns recognizable by EEG is very limited.

We conclude that while we have shown that BCI can be used as an interface for navigating VR in a CAVE-like setting, we are still far from being able to use the brain as a natural interface for such a purpose, let alone construct an interface which will be able to compete with traditional interface devices in accuracy and level of comfort. In order to reach such a goal, more research is needed that concentrates not only on BCI accuracy, but also on the overall participant experience.

Acknowledgments

This project was funded by the European Union Project PRESENCCIA, IST-2001-37927. We would also like to thank Marco Gillies and Vinoba Vinayagamoorthy for using the virtual street model. We are grateful for David Swapp for his support with the Cave.

References

- Bayliss, J. D. (2003). Use of the evoked potential P3 component for control in a virtual apartment. *IEEE Transactions Neural Systems Rehabilitation Engineering*, 11, 113–116.
- Bayliss, J. D. & Ballard, D. H. (2000). A virtual reality testbed for brain computer interface research. *IEEE Transactions of Rehabilitation Engineering*, 8(2), 188–190.
- Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. (1992). The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM*, 35(6), 65–72.
- Frecon, E., Smith, G., Steed, A., Stenius, M., & Stahl, O. (2001). An overview of the COVEN platform. *Presence: Teleoperators and Virtual Environments*, 10(1), 109–127.
- Garau, M., Ritter-Widenfeld, H., Antley, A., Friedman, D., Brogni, A., et al. (2004). Temporal and spatial variations in presence: A qualitative analysis. *7th Annual International Workshop on Presence*.
- Hettinger, L. J., & Riccio, G. E. (1992). Visually induced motion sickness in virtual environments. *Presence: Teleoperators and Virtual Environments*, 1(3), 306–310.
- Krausz, G., Scherer, R., Korisek, G., & Pfurtscheller, G. (2003). Critical decision-speed and information transfer rate in the “Graz brain-computer interface.” *Applied Psychophysiological Biofeedback*, 28(3), 233–240.
- Leeb, R., Scherer, R., Lee, F., Bischof, H., & Pfurtscheller, G. (2004). Navigation in virtual environments through motor imagery. 9th Computer Vision Winter Workshop, CVWW, '04, 99–108.
- Middendorf, M., McMillan, G., Calhoun, G., & Jones, K. S. (2000). Brain-computer interface based on the steady state visual evoked response. *IEEE Transactions on Rehabilitation Engineering*, 8(2), 211–214.
- Nelson, W., Hettinger, L., Cunningham, J., & Roe, M. (1997). Navigating through virtual flight environments us-

- ing brain-body-actuated control. *IEEE Virtual Reality Annual International Symposium*, 30–37.
- Pfurtscheller, G., & Neuper, C. (2001). Motor imagery and direct brain computer communication. *Proceedings of the IEEE*, 89(7), 1123–1134.
- Pfurtscheller, G., Neuper, C., Muller, G. R., Obermaier, B., Krausz, G., Schlögl, A., et al. (2003). Graz-BCI: State of the art and clinical applications. *IEEE Transactions of Neural Systems Rehabilitation Engineering*, 11(2), 177–180.
- Slater, M. (1999). Measuring presence: A response to the Witmer and Singer questionnaire. *Presence: Teleoperators and Virtual Environments*, 8(5), 560–566.
- Slater, M., & Usoh, M. (1993). Presence in immersive virtual environments. *IEEE Virtual Reality Annual International Symposium*, 33–40.
- Slater, M., Usoh, M., & Steed, A. (1995). Taking steps: The influence of a walking metaphor on presence in virtual reality. *ACM Transactions on Computer Human Interaction (TOCHI)*, 2(3), 201–219.
- Steed, A., Mortensen, J., & Frecon, E. (2001). Spelunking: Experiences using the DIVE system on CAVE-like platforms. *Immersive Projection Technologies and Virtual Environments*, 2001, 153–164.
- Usoh, M., Arthur, K., Whitton, M., Bastos, R., Steed, A., Slater, M., et al. (1999). Walking > walking-in-place > flying, in virtual environments. *SIGGRAPH 99*.
- Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller G., & Vaughan, T. M. (2002). Brain-computer interfaces for communication and control. *Clinical Neurophysiology*, 113(6), 767–791.

Appendix A: Subset of Presence Questionnaire: Overall Presence

A subset of questions related to overall presence from the SUT presence questionnaire (Slater, 1999). For all questions below the subjects were given choices from 1 to 7, and the answers were normalized such that higher scores indicate a higher sense of presence.

	L1 bar	L1 street	O3 bar	O3 street	S1 bar
Please rate your sense of being in the room, on the following scale from 1 to 7, where 7 represents your normal experience of being in a place.	5	6	3	2	4
To what extent were there times during the experience when the room was the reality for you and you almost forgot about the real world of the laboratory where the experience was really taking place?	4	4	1	1	2
When you think back about your experience, do you think of the room more as images you saw, or more as somewhere that you visited? (visited = 7)	3	6	1	1	2
During the course of the experience, which was strongest on the whole, your sense of being in the bar, or of being in the real world of the laboratory? (bar = 7)	6	5	1	2	5
During the time of the experience, did you often think to yourself that you were just standing in a laboratory or did the bar overwhelm you? (bar = 7)	5	4	1	1	4

Appendix B: Subset of Presence Questionnaire: Social Presence

The table below is an excerpt related to social presence from the SUT presence questionnaire

(Slater, 1999) for three subjects on two tasks. For all questions below the subjects were given choices from 1 to 7, and the answers were normalized such that higher scores indicate a higher sense of presence.

	L1 bar	L1 street	O3 bar	O3 street	S1 bar
During the course of the experience, did you have a sense that you were in the room with other people or did you have a sense of being alone?	5	4	7	1	2
How aware were you of the characters in the room?	5	5	6	6	3
How curious were you about the characters?	3	3	5	5	5
When you <i>first</i> saw the characters, to what extent did you respond to them as if they were real people?	5	3	1	1	4
Now consider your response over the course of the whole experience. To what extent did you respond to them as if they were real people?	4	1	1	1	4
To what extent did you have a sense of being in the same space as the characters?	6	7	1	5	4
To what extent did the presence of the characters affect the way you explored the space?	6	5	3	5	6
How much do you think you disturbed the characters in the room?	2	3	1	1	1
When you <i>first</i> saw them, did you respond to the characters more the way you would respond to people, or the way you would respond to a computer interface?	5	4	7	7	4
Now consider your response <i>over the course of the whole experience</i> . Did you respond to the characters more the way you would respond to people, or the way you would respond to a computer interface?	5	3	7	7	4