HUMAN–COMPUTER INTERACTION, 2010, Volume 25, pp. 67–93

ISSN: 0737-0024 print / 1532-7051 online
DOI: 10.1080/07370020903586688

Human–Computer Interface Issues in Controlling Virtual Reality With Brain–Computer Interface

Doron Friedman,[1] Robert Leeb,[2] Gert Pfurtscheller,[3] and Mel Slater[4]

[1]*The Interdisciplinary Center, Israel*
[2]*Graz University of Technology, Graz, Austria, and Ecole Polytechnique Fédérale de Lausanne*
[3]*Graz University of Technology, Austria*
[4]*ICREA-University of Barcelona and University College London*

We have integrated the Graz brain–computer interface (BCI) system with a highly immersive virtual reality (VR) Cave-like system. This setting allows for a new type of experience, whereby participants can control a virtual world using imagery of movement. However, current BCI systems still have many limitations. In this article we present two experiments exploring the different constraints posed by current BCI systems when used in VR. In the first experiment we let the participants make free choices during the experience and compare their BCI performance with participants using BCI without free choice; this is unlike most previous work in this area, in which participants are requested to obey cues. In the second experiment we allowed participants to control a virtual body with motor imagery. We provide both quantitative and subjective results, regarding both BCI accuracy and the nature of the subjective experience in this new type of setting.

Doron Friedman is a multidisciplinary scientist with an interest in virtual reality and artificial intelligence; he is a Lecturer in the Sammy Ofer School of Communications in the Interdisciplinary Center, Herzliya, Israel and an Honorary Lecturer in the Department of Computer Science, University College London. **Robert Leeb** is bio-medical engineer and computer scientist with an interest in brain–computer interfaces, bio-signal processing, rehabilitation engineering, and virtual reality systems; he is a Chair of Non-Invasive Brain-Machine Interface in the School of Engineering at the Ecole Polytechnique Fédérale de Lausanne in Switzerland. **Gert Pfurtscheller** is Full Professor of Medical Informatics and founding director of the Laboratory of Brain–Computer Interface at the Graz University of Technology in Austria; his research interests include functional brain topography using event-related desynchronization, the design of brain–computer communication systems, and navigation in virtual environments by a brain–computer interface. **Mel Slater** is a computer scientist with an interest in the application of virtual reality in psychology and neuroscience. He is an ICREA Research Professor in the Faculty of Psychology at the University of Barcelona, and Professor of Virtual Environments in the Department of Computer Science, University College London.

CONTENTS

1. INTRODUCTION

In 1965 Ivan Sutherland introduced the idea of the "Ultimate Display." The ripples of effect from his original concept and realization are still working themselves out today. In this research we explore one possible paradigm of the Ultimate Interface: One where people do not actually have to do anything physical to interact with and through a computer, but where the computer is directly attuned to their brain activity. A particularly exciting realm in which to investigate the possibilities inherent in this idea is within virtual reality (VR). Imagine a situation where a participant only needs to think about an action in order to make it happen: to move, to select an object, to shift gaze, to control the movements of their (virtual) body, or to design the environment itself, by "thought" alone. What would that be like?

There would be obvious practical advantages, such as for people with limited physical abilities, where much of the research behind this work originated (Dornhege, Millan, Hinterberger, McFarland, & Muller, 2007; Pfurtscheller & Neuper, 2001; Wolpaw, Birbaumer, McFarland, Pfurtscheller, & Vaughan, 2002).

Moreover, if successful, this line of research could lead to a paradigmatic revolution in the field of human–computer interaction (HCI), possibly a significant step following "direct manipulation interfaces" (Schneiderman, 1983)—where intention is mapped directly into interaction, rather than being conveyed through motor movements.

Such an interface can be made possible using a brain–computer interface (BCI). 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 motor actions,[1] from online electroencephalogram (EEG) signals. Such thought-related EEG changes are transformed into a control signal and associated to simple computer commands (i.e., cursor movement).

Highly immersive VR can be a safe and controlled replacement for using BCI in reality, and could thus serve as a transition phase from lab to real-world deployment. In addition, highly immersive VR may provide the subject with feedback that is most similar to the stimuli provided in the real world, which may improve adaptation to BCI and performance. Furthermore, participants find VR very enjoyable, which is critical for BCIs based on extensive training.

We have set up a system that connects the EEG-based Graz-BCI (Pfurtscheller & Neuper, 2001) to a highly immersive VR. The VR experience took place in a four-walled Cave-like system. A Cave is an approximately 3 m cubed area, with projection screens on the floor and three walls providing a stereo view. We use the term "Cave" in this article to refer to the generic type of system described in Cruz-Neira, Sandin, DeFanti, Kenyon, and Hart (1992). The BCI-Cave setting appears in Figure 1.

Using a BCI to control VR raises several major HCI issues: whether classification of "thought" patterns is continuous (asynchronous BCI) or only takes place in specific moments (synchronous BCI), the number of input classes recognized, the importance of feedback, and the nature of the mapping between imagery and resulting action in the virtual environment (VE). In this article we refer to these issues and present two case studies that specifically address two of these issues.

First, we needed to establish that it is possible to use BCI in a VR. Thus, it was necessary to accurately measure BCI success rate. The most reliable way of doing this is to give the subjects specific instructions about what they are supposed to "think." In a previous experiment (Friedman et al., 2007; Pfurtshceller et al., 2006) we have thus exposed participants in the Cave to two different auditory cues. One audio cue signaled the participants to activate one type of motor imagery, such as hand movement, and the other cue signaled another type of motor imagery, such as foot movement. By motor imagery, we refer to the participants imagining their hands, arms, or feet move, without actually moving them. As reported in the articles just mentioned, three participants were successful in performing two simple navigation tasks in the VR. However, in this previous experiment, the participants were not free to perform a task in the environment.

---

[1]The technical term used in the BCI community is *motor imagery*.

FIGURE 1. A participant in the virtual street, connected to the BCI equipment, inside the Cave-like system.

Our goal in the research described in this article is to move toward a scenario closer to real-world BCI usage. The first step taken here is to allow the participants to make their own choices, thus introducing "free will." The importance of free choice is recognized by BCI researchers; for example, Wolpaw et al. (2002) wrote, "Because they [most laboratory BCIs] need to measure communication speed and accuracy, laboratory BCIs usually tell their users what messages or commands to send. In real life the user picks the message" (p. 772).

In the first experiment reported here, 10 participants attempted to perform simple navigation tasks in a virtual street, using motor imagery of their right hand, left hand, or feet. Half of the participants were given a simple task and were asked to perform it with free choice. The second half of the participants performed a navigation task using our original BCI paradigm, which does not include free will. Section 3 provides details of the experiment and the results. While preparing for this experiment, we have evaluated five navigation paradigms, and we describe the lessons we have learned from these sessions. In addition, we present some qualitative data collected using questionnaires and interviews, regarding the subjective experience of navigating VR with BCI and free choice.

The second experiment considers the nature of visual feedback provided and its importance. A plausible initial hypothesis is that natural mapping between thought processes and feedback in the VE would improve the experience. This was explored in a second case study, concerning BCI control of a virtual body.

A key requirement for a successful experience in an immersive VE (IVE) is the representation of the participant, or its avatar (Pandzic, Thalmann, Capin, & Thalmann, 1997; Slater, Steed, McCarthy, & Maringelli, 1998; Slater, Usoh, & Steed, 1994). This article describes the first-ever study where participants control their own avatar using BCI. Three participants were able to use the Graz-BCI to control an avatar, and their subjective experience was assessed using questionnaires and a semistructured interview.

We are interested in the mapping between the thoughts used to activate the BCI and the resulting functionality in the VE, that is, if a person thinks about kicking a ball, the mapping would be natural if they would see their virtual foot kick a virtual ball, that would actually be kicked away. A one-to-one mapping seemingly makes intuitive sense, but having this mapping is constraining because we are limited in the scope of thought patterns that we can detect based on contemporary brain recording techniques. In addition, it precludes other more complex or more fanciful body image mappings; what if we want to experiment with lobster avatars?[2] Section 4 provides details of the experiment and the results.

2. BACKGROUND

2.1. Brain-Computer Interface

The possibility that people may be able to control computers by thought alone, based on real-time analysis of EEG waves, was already conceived as early as 1973 (Vidal, 1973). Recently, with advances in processing power, signal analysis, and neuro-scientific understanding of the brain, there is growing interest in BCI, and a few success stories. Current BCI research is focusing on developing a new communication alternative for patients with severe neuromuscular disorders, such as amyotrophic lateral sclerosis, brain stem stroke, and spinal cord injury (Wolpaw et al., 2002).

A BCI-system is, in general, composed of the following components: signal acquisition, preprocessing, feature extraction, classification (detection), and application interface. The signal acquisition component is responsible for recording the electrophysiological signals and providing the input to the BCI. Preprocessing includes artifact reduction (electrooculogram and electromyogram [EMG]), application of signal processing methods (i.e., low-pass or high-pass filters), methods to remove the influence of the line frequency, and in the case of multichannel data, the use of spatial filters (bipolar, Laplacian, common average reference).

After preprocessing, the signal is subjected to the feature extraction algorithm. The goal of this component is to find a suitable representation (signal features) of the electrophysiological data that simplifies the subsequent classification or detection of specific brain patterns. There is a variety of feature extraction methods used in BCI systems. A nonexhaustive list of these methods includes amplitude measures, band power, phase features, Hjorth parameters, autoregressive parameters, and wavelets. The task of the classifier component is to use the signal features provided by the feature extractor to assign the recorded samples of the signal to a category of brain patterns. In the most simple form, detection of a sin-

[2]See Jaron Lanier's "everybody can be a lobster" statement at http://www.edge.org/q2006/q06 7.html#lanier

gle brain pattern is sufficient, for instance, by means of a threshold method; more sophisticated classifications of different patterns depend on linear or nonlinear classifiers (Pfurtscheller & Neuper, 2001). The classifier output, which can be a simple on–off signal or a signal that encodes a number of different classes, is transformed into an appropriate signal that can then be used to control, for example, a VR system. Lotte, Congedo, Lecuyer, Lamarche, and Arnaldi (2007) provided an excellent review of classification algorithms used in BCI and provided guidelines for choosing the most suitable algorithm, based on BCI paradigm and EEG features.

General reviews of BCI include Wolpaw et al. (2002); Vallabhaneni, Wang, and He (2005), and Dornhege et al. (2007). BCI is still in its infancy. Most would argue that it has no advantage for healthy participants, and even use of BCI for patients, in clinical practice, is very limited (Kubler, Kotchoubey, Kaiser, & Wolpaw, 2001). Some companies have recently raised interest in BCI for the wide population, mainly entertainment.

Neuper, Scherer, Reiner, and Pfurtscheller (2005) studied the differential effects of kinesthetic and visual-motor mode of imagery in EEG. Classification accuracy for motor execution and observation of movement were larger than those for motor imagery. Kinesthetic motor imagery resulted in better classification than visual-motor imagery, and there were differences in the spatio-temporal patterns among those types of imagerys.

There are several methods for EEG-based BCI. Some are regulated by the BCI user, such as slow cortical potentials and sensorimotor rhythms, and some are elicited by visual, auditory, or tactile stimulation, mostly the P300 or steady-state evoked potentials. A description of the physiological basis of these methods can be found in Kubler and Muller (2007).

The different methods have advantages and drawbacks. The BCI community is usually concerned with accuracy, and sometimes with the throughput of the interface (bit-rate and number of commands available; (Schlogl, Kronegg, Huggins, & Mason, 2007). In this article we address issues that we believe to be important but are rarely addressed by the BCI community, mainly the type of mapping between "thought-patterns" and the functionality, and whether the BCI user is free to choose the BCI class.

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, Hettinger, Cunningham, and Roe (1997) were interested in BCI as a potential means for increasing the effectiveness of future tactical airborne crew stations. They investigated the usage of CyberLink™—an interface that uses a combination of EEG and 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-ft diameter dome display. Continuous feedback was provided by a graphical heads-up display. Participants were not given any BCI instructions or training. Task performance scores 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 Middendorf et al.'s experiment two methods were employed, and 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-sec period were higher than some threshold and rolled left if most of the samples were lower than another threshold. Most operators were able to reach 85 to 90% of success after 30 min of training.

Bayliss and Ballard (2000) used the P3 evoked potential (EP), a positive waveform occurring approximately 300 to 450 msec after an infrequent task-relevant stimulus. They used a head-mounted display (HMD)-based system. Participants 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 participants were driving a modified go-cart. Whenever a red light was displayed, data were recorded continuously from –100 to 1,000 msec. 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 IVE. Participants 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 a HMD. Using BCI, participants could switch the objects on and off or cause the animated character to appear or disappear. The BCI worked as follows: approximately once per second a semitransparent sphere would appear on a randomly selected object, for 250 msec. Participants were asked to count the flashes on a specific object (to make the stimulus task-related, as P3 requires). An epoch size from –100 msec (before the stimulus) to 1,500 msec was specified. Text instructions in the bottom of the screen indicated the goal object. The participant had to count the flashes for that object only. The participant was given a visual feedback when a goal was achieved, that is, when a P3 event was recorded when the target object was flashing. Participants were able to achieve approximately three goals per minute. Bayliss found no significant difference in BCI performance between IVE and a computer monitor. Most participants preferred the IVE environment; all of them liked the fixed-display condition (looking through a fixed HMD) the least.

This previous research into VR and BCI was all based on several types of evoked responses. Our research is based on a different BCI paradigm that exploits motor imagery. Motor imagery may be seen as mental rehearsal of a motor act without any overt motor output. It is broadly accepted that mental imagination of movements involves similar brain regions that are involved in programming and preparing such movements (Jeannerod & Frak, 1999). According to this view, the main difference between motor performance and motor imagery is that in the latter case execution would be blocked at some corticospinal level. Functional brain imaging studies monitoring changes in the metabolism revealed, indeed, similar activation patterns during motor imagery and actual movement performance (Lotze et al., 1999).

Motor imagery has been shown to represent an efficient mental strategy to operate a BCI (Pfurtscheller & Neuper, 2001). The imagination of movement of different body parts (e.g., right hand, left hand, foot or tongue movement) results in a characteristic change of the EEG over the sensorimotor cortex of a participant. Typically we allow each participant, during training, to find the specific type of movement that would work best for him or her, for example, for foot movement one participant may find that imagining cycling works best, whereas another would imagine kicking a ball.

BCI is typically considered as a means of controlling a device or an application. This scenario may be useful for populations with severe disabilities, but at the current state of BCI has limited benefits for healthy people. We note that there may be other scenarios, in which real-time analysis of EEG is only one out of several communication channels (e.g., an application that dynamically adapts to the participant's cognitive load). Such a paradigm may be useful for healthy participants as well; however, these are outside the scope of the scope of this article.

The main achievements of BCI to date are discussed in Kubler and Muller (2007). Noninvasive event-related (such as P300- and SSVEP-based) BCIs have achieved relatively high transfer rates, and tasks such as spelling or two-dimensional cursor control were successfully carried out. One of the main breakthroughs of noninvasive BCIs would be if "dry" electrodes are developed: electrodes that do not require conductive gel and still provide clean signals. If such electrodes are made available, we foresee much wider adoption of BCI.

2.2. Virtual Reality

HCI, specifically VR research, is continuously striving toward natural and seamless human–computer interfaces, and the existing interfaces for locomotion through 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 are 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 et al. (1995) 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 (for recent reviews of the concept of presence, see Vives & Slater, 2005, and Riva, Davide, & IJsselsteijn, 2003). In a later experiment (Usoh et al., 1999), walking in place was compared with real walking, and in terms of the reported sense of presence, the results were not much different. One of our questions in the line of research described here is, rather than walking in place, what would it be like if we were able to navigate a VE by merely imagining ourselves walking?

2.3. Brain-Computer Interface: HCI Issues

In previous experiments (Friedman et al., 2007; Leeb et al., 2006; Pfurtshceller et al., 2006), we have allowed participants to navigate a virtual street using BCI in a Cave-like (Cruz-Neira et al., 1992) system. Our results in that previous experiment provided some evidence that a highly immersive environment such as a Cave may not only improve user motivation but may also facilitate BCI accuracy. This suggests that there is a great potential in using VR with BCI. However, our research has also made us aware of the many limitations and design issues that come into play when using BCI as an interface for control in VR, which we now consider.

The first issue is the number of different events (or classes) distinguished in real time, through the analysis of EEG. As we add more classes, accuracy quickly drops, and the number of EEG channels (recorded brain areas) needs to grow, which makes the sessions more complex and time consuming.

Another limitation is that BCI is often synchronous, or trigger-based, that is, the classification is not applied continuously, but only in specific time windows following an external cue, such as a short sound after which participants are required to have the appropriate "thoughts." Naturally, such a trigger-based interface is limited in comparison to most traditional input devices. Asynchronous BCI is possible; the first asynchronous BCI was demonstrated by Mason and Birch (2000), but accuracy is compromised (Millan & Mourino, 2003). We have applied asynchronous BCI in the Cave, as reported elsewhere (Leeb, Friedman, et al., 2007).

Brain states for BCI can be achieved by two methods. The first is self-regulation of the brain activity, which is achieved through operant conditioning (Dragoi & Staddon, 1999). Another possibility is for the participants to perform different mental tasks (motor imagery, mental rotation, mental arithmetic) and use pattern recognition methods to detect the different thoughts. The Graz-BCI uses a combination of these two approaches: It uses a defined mental strategy (i.e., motor imagery) but also provides continuous feedback (i.e., feedback learning). It is clear from the literature (e.g., Utz, 1994) that both search for a strategy and reinforcement affect the result. However, when BCI is eventually deployed to perform a task (either in a virtual world or in the real world), it cannot rely on the effect of conditioning, because users need to be able to make their own decisions. Thus, free choice is of high interest and practical implications, and this was addressed by the first experiment.

Wolpaw et al. (2002) highlighted the importance of feedback for BCI. To be effective, the feedback needs to be immediate. However, providing continuous and immediate feedback causes a problem. If we look at the accuracy of classification over time, we see that there is typically a delay of 1 to 3 sec between the onset of the trigger and the optimal classification. The typical approach, which we also adopt here, is to provide feedback for the whole classification duration (approximately 4 sec), even though we know the classification is rarely correct throughout this whole duration. Figure 2 shows the data from a typical training run: In this case the error level drops to optimum only 2-3+ sec after the trigger, and then rises again.

FIGURE 2: (a) A cross is displayed from time 0, and an arrow cue is given at second 3 for a duration of 1.25 sec, which indicated to the participant what they should "think." The cross appears for 4.25 sec after the display of the arrow and indicates the time during which the participant has to focus on the imagery. The cross and the arrow disappear simultaneously. (b) Classification error over time, averaged more than 40 triggers in one run—example of 1 participant.

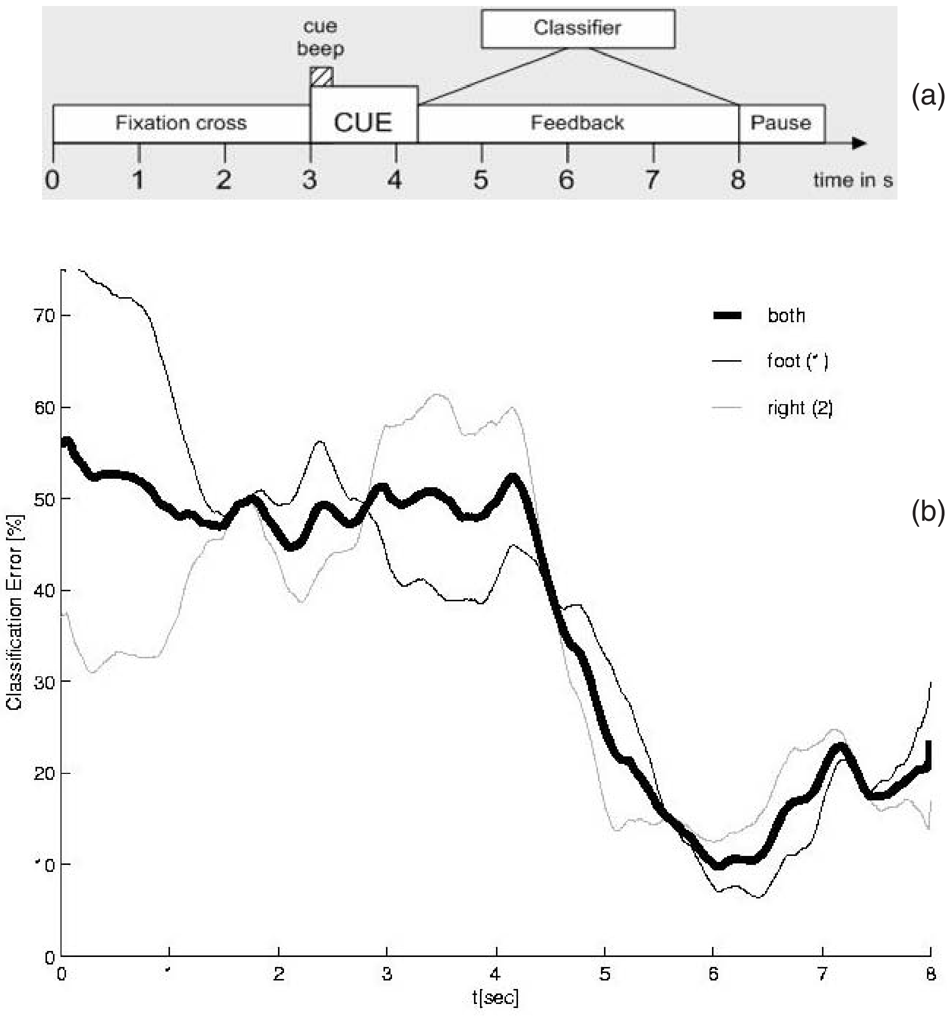

Another issue is the importance of an intuitive mapping between the mental activity used for control and the visual feedback; this is addressed by the second experiment, in Section 4. In general, it is very difficult to study the exact relationship between EEG patterns and thoughts; in a sense, this may entail solving the mind–body problem, which is considered by many philosophers to be the major philosophical problem in western philosophy. Clearly, the same EEG signal may be related with many thoughts, and the same thought may result in different EEG patterns, even for the same person.

3. EXPERIMENT 1: BCI-VR NAVIGATION WITH FREE CHOICE

3.1. Method

Participants and Training

The study was performed on 10 healthy volunteers aged 20 to 38 years (*M* age = 27.7 years), 4 female and 6 male. All participants had no history of neurological disease, gave formal consent to participate in the study, and were paid 7 GBP per hour. Each session lasted 2 to 4 hr, and each participant went through one session.

BCI experiments with the Graz BCI paradigm are usually carried out after extensive training, typically lasting a few weeks. As part of our goal in this research, to investigate BCI-VR in a less restricting context, we were interested in minimal training.

Our procedure was based on screening subjects for BCI performance using the traditional, monitor-based BCI. When participants reached an acceptable level, they were immediately moved to the VR Cave system. We arbitrarily decided to set a BCI success rate above 75% as satisfactory. Based on previous experience with the Graz-BCI, we expected a significant percentage of a random population to be able to achieve this level of BCI success (Guger, Edlinger, Harkam, Niedermayer, & Pfurtscheller, 2003). In our case, we were able to keep 10 out of 14 participants.

Each participant first took part in two to three training runs without feedback. In each run the participant had to imagine a movement of both their legs or a movement of their right hand in response to a visual cue-stimulus presented on a computer monitor, in the form of an arrow pointing downwards or to the right, respectively (Figure 3). In addition to the visual cue an auditory cue stimulus was also given either as a single beep (hand imagery) or as a double beep (leg imagery). Each trial started with a fixation cross (Second 0) followed at Second 3 by the cue-stimulus presented for 1.25 sec. There was a random duration interval in the range from 0.5 to 2 sec between the trials (see Figure 2).

Forty EEG trials, 20 for every class, were recorded in one run. The EEG trials from runs without feedback were used to set up a classifier for discriminating between the two different mental states. In further runs, visual feedback in the form of a moving bar was given to inform the participant about the accuracy of the classification during each imagery task (i.e., classification of imagined right-hand movement was represented by the bar moving to the right, classification of imagined foot movement by the bar move vertically; see Figure 3).

This allowed us to have the training and the experiment on the same day, which is important, as the classifier tends to change over days; obtaining a stable, long-term classifier is much more difficult.

EEG Recording

Three EEG channels were recorded bipolarly (two electrodes for each channel). Electrodes were placed 2.5 cm anterior and 2.5 cm posterior to positions C3, C4, and Cz

FIGURE 3. "Traditional" BCI in front of a monitor: the arrow on the red cross indicates to the participants whether they should imagine moving their hand or their feet. Participants need to keep concentrating on this thought as long as the arrow is displayed, for 4.25 sec. If a classifier is already available, the white bar provides immediate and continuous feedback for 4.25 sec.

of the "international 10-20 system," which is a standard for electrode placement based on the location of the cerebral cortical regions. The EEG was amplified between 0.5 and 30 Hz by an EEG amplifier (g.tec Guger Technologies, Graz, Austria) and processed in real time. Sampling frequency was 250 Hz. Note that we have used a paradigm that was well established for the Graz BCI, hence using more electrodes was not necessary.

We instructed the participants not to move their arms, and this was verified visually during the experiments. In previous work, Leeb, Lee, et al. (2007) recorded the EMG on both forearms, in addition to the EEG. During the periods of imagery they found no significant EMG activity and showed they could not use the EMG for classification. Another way to verify there were no EMG artifacts is by observing the EEG spectrum; any "contamination" is easily detected, as EMG contains higher frequency components.

Feature Extraction and Classification

BCI systems apply an online classification to the EEG signal. Two frequency bands selected from each EEG channel served as features for the algorithm: the logarithmic band power was calculated in the alpha (8–12 Hz) and beta (16–24 Hz) bands, over 1-sec epochs. These features were classified with Fisher's linear discriminant analysis and transformed into a binary control signal.

Virtual Reality

The experiments were carried out in a four-sided Cave-hybrid (Cruz-Neira et al., 1992) system. The particular system used was a Trimension ReaCTor, with an Intersense IS900 head-tracking. The applications were implemented on top of the DIVE software (Frecon, Smith, Steed, Stenius, & Stahl, 2001; Steed, Mortensen, & Frecon, 2001).

We have used a VE depicting a virtual high street, with shops on both sides, and populated by 16 virtual characters (see Figure 1). The participant's task was to navigate and reach the characters.

A communication system called Virtual Reality Peripheral Network (http://www.cs.unc.edu/Research/vrpn/) was used for the communication between the PC running the BCI and the VR host. A diagram of the integrated system appears in Figure 4. More details about the BCI-VR integrated system are provided in Friedman et al. (2007).

The Experimental Conditions

We wanted to compare two conditions: one in which we use the previous paradigm (Friedman et al., 2007; Leeb et al., 2006), where the participants are instructed what to think, and one where the participants are free to choose the type of motor imagery. Both cases were comprised of a 1-dimensional navigation task. In the original condition, the participants were instructed to follow the cues. In the new condition reported here, the participants were given an overall goal that they had to follow, rather than specific cues.

The Experimental Condition: Go and Touch. We wanted the task to be realistic for a VE scenario yet possible with the major limitations imposed by the BCI.

FIGURE 4. The Graz-BCI connected to the UCL VR Cave.

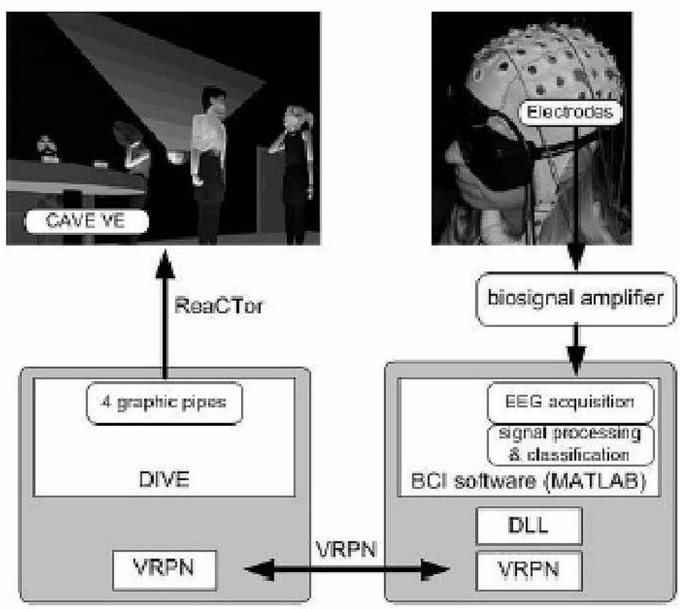

We tested several variants of tasks until we came up with a paradigm that is appropriate for comparison. The lessons learned from evaluating various paradigms are detailed in the Other Paradigms and Tasks section.

Ideally, the BCI would be asynchronous. Such asynchronous BCI is, as previously mentioned, extremely difficult to achieve, and we addressed this in another VR-BCI experiment (Leeb, Friedman, et al., 2007). In this experiment, we use synchronous (trigger-based) BCI, but in a way that still allows free will, rather than instructing the participant specifically what to think for each trigger.

The free-choice paradigm we tested involves moving forward to reach objects and then touching them. This is a generic paradigm for exploring and acting within a VE. In our specific case we used the "sleeping beauty" scenario: Participants used foot imagery and head rotation to move forward toward the characters. Within a certain proximity from a character, the participant had to imagine a hand movement. There was no visual feedback of hand motion. When characters were "touched" in this way they would "wake up" and start walking. This mapping between thought patterns and VE functionality is thus very natural: walking using foot imagery and acting on an object using hand imagery.

Note that unlike in the typical paradigm, in the free-choice paradigm we use only one type of audio trigger, upon which the participants could always use one of two motor imagery patterns. For each audio trigger the VR system receives up to 80 updates over an epoch of 4.16 sec. These updates are used to control the VE and are used to provide the participants with immediate feedback. Note that the audio triggers are very short. Thus, in our VR conditions, the participant does not have precise feedback indicating when the epoch ends; this is unlike the desktop-based BCI, where the termination of the epoch is indicated by the disappearance of the cross.

The VE functionality was based on each one of the 80 classification results per trigger. For each foot classification, the participant moved forward in the environment by a constant small distance. For participants to successfully touch a virtual character, they had to be within a touch range from a character. Then, given a trigger, they had to imagine right-hand movement. The virtual character was considered to be touched only if the majority of the updates for that trigger indicated right-hand imagery.

In this paradigm the direction of movement is determined by the participant's gaze direction; this is possible because the participant is head-tracked in the Cave. Thus, the paradigm allows two-dimensional navigation. Because head motion might interfere with BCI classification, participants were asked to move their heads only in the time gap between two BCI epochs. The gap between two BCI epochs is a random duration between 2 and 4 sec; the randomness is introduced to avoid habituation.

The Control Condition. The control condition, in which the participants were instructed what to think, was repeated with the same paradigm as reported in our previous experiment (Friedman et al., 2007; Leeb et al., 2006). We tested new participants, as we needed data from participants that went through the same minimum training.

Using this paradigm the participant had no free choice. 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, he or she had to imagine right-hand movement. If the classification indicated hand imagery when ordered to walk, the participant would stay in place. If the classification indicated foot imagery when ordered to stop, the participant would go backward. This "punishment" was introduced so that the participants will not be able to adopt a strategy of always imagining their foot.

3.2. Results

Evaluating Free Choice

Our main question was whether there is a significant difference between a free-choice paradigm as compared to an instructed BCI paradigm. Each participant carried out two sessions. Each session includes 40 triggers, with 80 classification results per trigger.

Because the task was very simple, we assume that we can completely predict what the participant was trying to think for each of the triggers—if the participant was in front of a character we could assume they were attempting hand imagery, whereas otherwise we assume they would imagine their feet. Although this cannot be assumed beyond doubt, we found this assumption to be plausible based on postexperimental debriefing, that is, the participants reported they tried to perform the task correctly and that they found the cognitive part of the task to be trivial.

The BCI accuracy of the 10 participants in all the different runs under both conditions ranged from 65% to 95%. The mean BCI accuracy in the free-choice condition was 75.0% ($\sigma = 7.9$, $n = 10$), lower than the control condition, which was 82.1% ($\sigma = 6.7$, $n = 10$). In this case, even though there were large intersubject differences, the difference between the two conditions was found to be significant with a two-tailed t test ($p = .04$).

In general, we suggest a normalization of BCI performance data based on the monitor training phase, which could be used as a baseline, that is, rather than comparing absolute BCI performance of participants, we can compare their BCI performance in the experiment relative to their BCI performance in, say, the last two training stages. This would, ideally, help in overcoming the large interpersonal differences in BCI performance. In the experimented reported here the difference between the two conditions was highly significant, so there was no need for such normalization.

Subjective Results

The control of a virtual environment using BCI is a new type of experience, and we were interested in getting some insight into the subjective experiences of the participants. We thus used a combination of questionnaires and semistructured interviews. The goal of the subjective questionnaires and interviews is exploratory. We hope to partially reconstruct the subjective experience to gain insight into this novel experience; this is a type of ideographic study (Kelly & Main, 1978), which is part of our research on presence in VR; additional details on our methodology can be found in

Friedman et al. (2006), and for recent reviews of the concept of presence, see Vives and Slater (2005) and Riva et al. (2003).

Subjective impressions of people, unlike their BCI accuracy, are dependent on contingent factors such as social background, video game exposure, and so on. Next we describe what our participants reported, but there is no way that this can be extrapolated, and a study with a larger number of participants is necessary.

The semistructured interview included 10 questions. The interviews were tape-recorded and analyzed. Here we report a few themes that came up.

Participant M3 reported a high sense of presence. He mentioned, "I forgot the surrounding environment of the lab. Every time I moved forward I felt the environment being an extension of myself." He later said, "I felt like I was located in a street in America." Even participant M2, who reported low presence, remarked, "… but the times where the environment became more of a dominant reality is when I was trying to move forward."

When comparing the monitor-based BCI to the Cave BCI, the typical response was similar to the following response made by participant M3: "In the Cave it was more fun because you were in an environment that wasn't really there. But that also means more distraction." Participant M1 remarked, "The task was easier in the VR but only with thinking about the feet because it results in something happening." All participants who had even partial success mentioned that moving forward by imagining foot movement is a very enjoyable experience. This is not surprising but may be important in its own right: BCI training is typically very long and exhausting; VR may prove useful in significantly improving the training experience and increase motivation.

Participants were asked whether they felt they controlled the environment. Though difficult to quantify, it seemed that, ironically, the participants that experienced the instructed-BCI condition reported a higher level of control. This could be due to the fact that they were, on average, more successful in the BCI task.

Participants were not specifically asked about free choice, but 3 of the 5 participants in the experimental condition referred to an effect of conditioning. They reported that after going through this training, it was difficult for them to use the "free-choice" BCI: they were expecting the trigger to "tell them what to do."

Participant F2 reported that, unlike in the BCI training, she had a very clear goal (i.e., reaching the virtual characters). For many of the triggers, she said her thoughts were focused on reaching the target, rather than on her feet. In this way VR could be an obstacle, because the BCI is tuned to pick up specific thought patterns, and not the thought patterns involved in obtaining a goal. For example, during adapting to the traditional BCI participants may find out that imagining a pedaling motion with their feet works best. A context of walking in a street may divert them from this specific motor image, which might impede their BCI performance.

Other Paradigms and Tasks

While preparing for the experiment reported here we evaluated different navigation tasks in the same VE. In addition to the paradigm that was eventually selected (as described in The Experimental Conditions section), we tested 8 participants overall (7 male, 1 female, average age of 28 years), using three more paradigms. Each participant typically carried out one to three sessions, with 40 triggers per session. Each participant experienced only one Cave condition, and none of the participants described in this section participated in the main experiment previously reported. By investigating various tasks and BCI paradigms we gained some more insights about how to use BCI as an interface in VR; these are detailed in the rest of this section.

Speed Control. In this paradigm the participant uses foot or right-hand imagery to accelerate and left-hand imagery to decelerate. There is a maximum speed and the speed cannot go below zero, that is, the participants were theoretically able to reach full stop but could not move backward. The motion direction was dictated by the participant's gaze direction as measured by the head tracker. Thus, the participant could navigate in two dimensions. Note that the mapping in this task between the thought patterns and the resulting functionality is not very intuitive. The task was to walk down the virtual street and look for a specific character inside one of the virtual shops.

We have tested this paradigm with two participants, and three runs were valid. The participants had 77% and 80% BCI success in the training phase prior to being introduced to the Cave. We tried different calibrations of the acceleration and deceleration speeds, but in all cases the participants were not able to carry out the task properly; movement inside the shops was too fast to allow them to look around. If we calibrated the speed to be slower, the rate was too slow to move forward down the street to get to the shops in the first place. This result is not specific to BCI control but is related to a well-known problem in VE navigation. Mackinlay, Card, and Robertson (1990) observed that most navigation techniques use high velocities to cover distances rapidly, but high velocities are hard to control near objects. The solution they suggest does not lend itself naturally to the Graz-BCI, as it requires specifying the target navigation point. We could still imagine context-dependent calibration of speed; for example, the system may be able to know whether the participant is in a small, confined space or in a large space, and automatically calibrate the distance units accordingly. We have not attempted this direction.

Sideways Control. Using this paradigm the participant was moving forward in constant speed in the virtual street. This was the only paradigm in which gaze control was not used. For each trigger, the participant could think about her left or right hand and move left or right, respectively. The task was to "run over" the virtual characters standing in the virtual street.

We tried this paradigm only with 1 subject, who had 87% classification success that day. The participant was able to touch some of the characters: 2, 1, and 3 characters in three consecutive runs. The participant found the task relatively easy and intuitive, even though he missed a few characters in each run.

Despite the success, this paradigm has a few limitations; for example, it is not possible for the participant to stop, slow down, or even to move straight ahead. Thus, we did not proceed with this paradigm.

Navigating Forward and Backward. Using this paradigm the participants use one type of motor imagery to move forward and another to move backward.

Typically they would imagine their feet to move forward and the right hand to move backward. In some cases we tried to allow the participant to move forward using right-hand imagery and backward using left-hand imagery. The task was very simple: The participants were asked to move forward in the street and collide with a virtual character wearing orange clothes. This character was not visible from the start point and required the participants to move forward before even seeing the character. When they reached the target, they were asked to go backward until they reached the entry point. The direction of motion was dictated by the participant's gaze direction, so this was a two-dimensional navigation task, but the BCI only controlled one dimension.

With this paradigm 3 participants performed seven valid runs. The participants had BCI success rates in the training phase varying between 77% and 95%. We are able to normalize the Cave results as follows: For the participants, we split the data into two parts: one in which they had to move forward, and the other after they reached the target and had to go backward.

The most successful participant had an average of 64% success moving forward ($\sigma = 9.8$, $n = 4$) and 62% backward ($\sigma = 15.2$, $n = 4$) over the four runs; this is short to the participant's BCI training performance, which was 77%. There seems to be no significant difference between moving forward and backward for that participant.

Of interest, a participant who had 95% success in the training phase was not successful in this task in the Cave. Rather than using foot imagery to move forward he used hand imagery. In the interview he mentioned that navigation forward and backward using right-hand and left-hand imagery is very counterintuitive and confusing. Additional data are needed to establish if the type of imagery (foot vs. hands) is responsible for the difference in performance in the Cave.

4. EXPERIMENT 2: CONTROLLING A VIRTUAL BODY WITH BCI

We have used the same BCI-Cave setting as described in Section 3; in this section we only mention the differences and detail the results.

4.1. Method

Participants and Training

Eleven participants went through "traditional" (2D, monitor-based) BCI training, and the top three were selected for the actual IVE study. It is known that a small percentage of the population can easily adapt to the BCI and a larger majority can reach similar accuracy levels, but only with long periods of training (Guger et al., 2003), thus typically 2 to 5 participants are used to prove the feasibility of a system. Because we were also interested in comparing between two conditions, we had each participant repeat each condition four times.

Each participant first took part in two or three training runs without feedback, as described in the Experiment 1 Participants and Training section. Two participants had a high BCI success rate in their first few runs: One had more than 85% accuracy and the other more than 90% accuracy. Finding a third participant was more difficult. After a few weeks of training, 3 other participants reached approximately 75% accuracy, showing improvement over time, but the improvement was slow.

Eventually, the study proceeded with 3 participants: 2 female and 1 male (aged 21, 49, and 27, respectively). All participants were right handed, without a history of neurological disease. Participants gave formal consent to participate in the study and were paid 7 GBP per hour. Each session lasted 3 to 4 hr, and each participant went through two sessions.

The Virtual Environment

The environment included two furnished virtual rooms. The avatar was projected (using stereo display) to appear standing approximately half a meter in front of the participant, who was sitting on a chair. The avatars were matched for gender with the participant (see Figure 5).

Research with own body representation (or avatars) may be more naturally carried out with an HMD, in which the participants cannot see their own body. However, in previous research we found that BCI participants prefer the Cave over the HMD: they feel more comfortable for longer durations (Friedman et al., 2007).

Experimental Conditions

The visual feedback was different in two conditions. In the first condition, which we call the normal condition, the mapping between the thought pattern and result in the IVE was intuitive: When the participants imagined moving their right arm the avatar would wave its right arm, and when they imagined moving their feet the avatar would start walking forward slowly. In the second condition the mapping was reversed: When the participants imagined moving their right arm the avatar would start walking, and when the participants imagined moving their feet the avatar would wave its arm. The feedback was continuous for the same duration as in the monitor-based BCI training (4.25 sec). In both conditions, the audio triggers were the same as in the training phase: A single beep indicated that the participants need to think about their arm, and a double beep indicated they need to think about their legs.

4.2. Results

BCI Accuracy

Each participant carried out four runs of both conditions, thus eight runs in total. Each run included 40 trigger events, and each trigger was followed by 80 classification results, one every approximately 50 msec. Thus, the data include eight runs per

FIGURE 5. (a) A female participant and her avatar in the virtual room. The participant is connected to the BCI equipment, inside the Cave-like system. (b) A male participant with the male avatar, in the same setting.

(a)

(b)

participant, and each run includes 3,200 trials. BCI accuracy is determined by the percentage of successful trials.

To test the significance of the results we carried out the equivalent of a two-way analysis of variance, using the number of successes out of the 12,800 trials in each of the conditions. In this analysis the response variable is therefore taken as a binomial distribution (rather than Gaussian) and it is a standard application of logistic regression. The results show that there were highly significant differences between the 3 participants (at a significance level that is too small to be quoted). Participant M1 had the highest success rate (94%), participant F1 had the next highest (86%), and participant F2 the lowest (81%)—and these are in keeping with what is typically found in BCI experiments. The raw figures show that in the normal condition the success rate was 86.7% and in the reverse condition was 87.7%, and with $n = 12{,}800$ per condition, this difference is significant. However, this does not take into account the differences between the participants—because the very large advantage of the reverse condition for participant F1 (88% reverse compared to 84% normal) distorts the overall result. For participant M1 the reverse condition is significantly higher than the normal condition ($z = -11.3$, $p = 0$) for participant F2 there is no significant difference between the reverse and normal condition ($z = 1.02$, $p = .31$) and for participant F1 the normal condition is significantly higher than the reverse condition ($z = 3.88$, $p = 1.0e\text{-}4$). These are carried out using a normal test for the difference between proportions. Thus, overall, no particular conclusion can be drawn one way or another about the effectiveness of the mapping in terms of BCI performance. Figure 6 depicts the performance of the three participants in the two conditions.

Subjective Results

The control of a virtual body using BCI is a completely new type of experience, and we were interested in getting some insight into the subjective experiences of the participants. We thus used a combination of questionnaires and semistructured interviews. The goal of the subjective questionnaires and interviews is exploratory. We hope to partially reconstruct the subjective experience to gain insight into this novel experience; this is a type of ideographic study (Kelly & Main, 1978).

After the first IVE session, each participant completed several questionnaires: the SUS presence questionnaire (Slater et al., 1994), the Trinity questionnaire for body plasticity (Desmond, Horgan, & MacLachlan, 2002), and a questionnaire regarding body projection: When a person has the sensation that an object (whether real or virtual) is experienced as part of his or her own body, this is referred to as "body projection." The most famous example of this is the rubber arm illusion (Armel & Ramachandran, 2003; Botvinick & Cohen, 1998). To evaluate whether this type of body projection was experienced by our participants, we have also administered a questionnaire recently designed in the UCL VECG lab for that purpose.

The questionnaires are comprised of 7-point or 5-point Likert-scale questions. First, all questions were normalized so that all low and high rates indicate the same

FIGURE 6. BCI error levels of the three participants in the two experimental conditions.

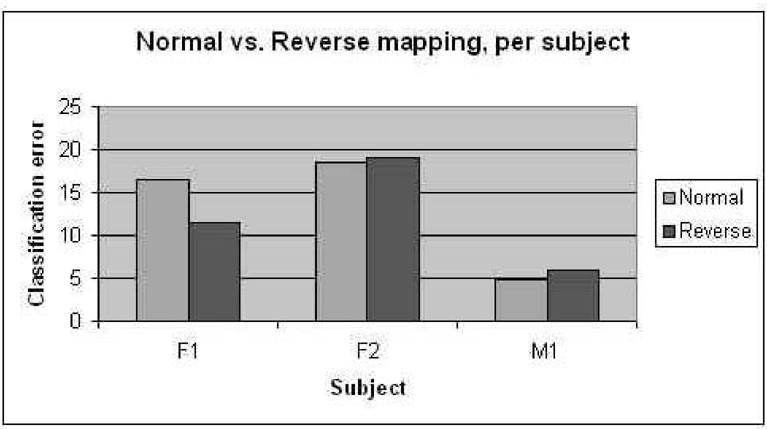

trend, for example, low presence would always correspond to a low rating. Then we counted how many extreme (very low or very high) answers each participant provides. (For 7-point questions, 1 and 2 were considered low and 6 and 7 were considered high, and for 5-point questions only 1 and 5 were considered extreme.) By subtracting the number of high scores from the number of low scores, we can classify the result of that questionnaire into three categories: low, high, or neutral. Our three participants showed consistency in their answers; there was no case where there were both high and low scores for the same questionnaire. Figure 7 summarizes the results, which were also used to complement the interviews in gaining an insight into the participant's experience.

After completing the questionnaires, the participants went through a semi-structured interview. The interviews were audio-taped and transcribed. Such interview agendas are designed in advance to identify logically ordered themes, following the recommendations of Smith (1977). We asked open-ended questions, and intervention was minimized to occasional neutral questions to encourage the participants to continue.

IVE-BCI versus Traditional BCI. Participants F1 and F2 thought the IVE-based BCI was easier (although they did not actually improve their BCI performance). Participant F1 compared the monitor-based BCI (which she refers to as a "little line") with the IVE experience: "I felt it was easier to make her do things. Because something was actually happening. Because when you're thinking about your feet but it's just a little line whereas if you're thinking about your feet and she moves it's, I don't why, it just seemed make to more sense." Participant M1, who reported very low presence, mentioned the IVE was more enjoyable.

Sense of Presence. One of the participants (F2) reported a high level of presence both in the questionnaire and in the interview. She related that to improved BCI performance: "At moments I had time to look around. And actually, then I really started—it became easy—walking, moving. … It felt in a strange way that everything became faster; time felt different. I left the other thoughts. It was a very different experience. It wasn't focused on the task—I was moving. I wasn't aware of doing the task. I was less aware of the signals, more aware of the environment. Less aware of you somewhere behind. Felt less as a task. I had the feeling: 'let me loose here.' I could have been able to do other things. It felt like a possibility, a reality."

Relationship With Avatar. Note that the participants were not told that the virtual body is intended to be their avatar, and in principle there is no reason why people should associate this virtual body with themselves. However, 2 of the 3 participants (M1

FIGURE 7. Summary of Questionnaire Results. Result in each category can be either high (+), low (–), or average (0).

Participant	Plasticity	Body Projection	Presence
1F	+	+	0
1M	–	–	–
2F	–	0	+

and F2) referred to the virtual body as a puppet controlled by them, which is a typical way to regard an avatar. The third participant (F1) even occasionally referred to it in first person.

Participant F1 seemed to have the highest level of projection of her body to the avatar. This was evident not only from the questionnaire but also during the sessions. At first, the participant referred to the avatar as "she," but after a few runs she started referring to it as "I." In the questionnaires this participant reported a medium level of presence and a high degree of body plasticity. In the interview, this participant said, "Although I was controlling her, I wasn't moving my hand. And I'd know if I was moving my hand." However, later she added, " … Oh yeah. It's because I, my brain, did move the hand. Towards the end I did feel it was representing me. I always felt like it was representing me but I didn't feel it was a part of me. … It's difficult. When you think about moving your hands you know whether you're moving your hands or not. If she was moving her hand mine wasn't moving. So she can't really be a part of me. Cause to feel the hand moving you'd have to feel the air going past it. But the more you were in the more comfortable you would become with that becoming you. It would just be a different type of you. Like a different version of you, almost. But it will never be you. … First like another body. Most of the time."

Participant F2 reported higher presence but lower levels of body plasticity and body projection. "I couldn't think of her as myself. I was trying to get into her skin, it was frustrating when I couldn't. When I was successful I was becoming closer, I was becoming her. Or she was becoming me. I'm still saying that for me to experience my movement somehow she was a distraction. Thinking of movement I could have done better without her. We didn't click. … But the connection was more like a puppet master. … Get rid of her. Just let me move in the environment—that was amazing. She was the task I was supposed to do."

Participant M1 reported low presence and low body projection. In the interview, he said, "First I thought it was another person standing in front of me. I thought what the objective was. I was wondering what would happen to this person. I didn't feel as it being my body but I felt I had some control of the person or of the body standing in front of me. … I would best describe it like a puppet."

Mapping of Imagery Type to Avatar Motion. Note that although we consider the mapping between type of motor imagery and resulting avatar action to be natural in the normal condition, the mapping is not necessarily perfect. For example, participants were not instructed to think of a particular leg motion, and thus they imagined cycling or kicking, whereas the avatar, although it moved its feet, would perform a different action: walking. This could be confusing for the participants, and might even divert them from the specific imagery they have been trained with. In the interviews, all participants replied that this mismatch was not a problem and that the feedback seemed appropriate. We do not know if this mismatch affects the BCI performance.

Participants F1 and F2 (who experienced medium and high presence, respectively) mentioned that the fact that they were sitting and the avatar was standing was more problematic. F2: "I usually do not walk forward while I am sitting down." We, of course, anticipated this problem. In pilot runs we tried to have the subject stand in the

Cave; this proved uncomfortable and generated too many motion artifacts in the EEG signals. We could have had the avatar sit down, but that was not the point of the study; in our vision for a future you project your body onto the avatar's body, and then the avatar can be free to operate in the VE, controlled by your "thoughts."

When asked about the difference between the two conditions, all participants mentioned that they had to concentrate more in the reverse condition. This may be an explanation as to why they sometimes performed better in the reverse condition than in the normal condition. F1: "It was confusing, but I didn't find it difficult." F2: "I'm not sure if seeing the feedback was so confusing—just adding all these layers. It made it difficult to concentrate." M1 reported very low presence, yet mentioned, "I was surprised that when it was reversed I found it harder to concentrate. It made me confused. At all times, not only in the beginning. So there must have been something on another level—I must have been influenced."

5. DISCUSSION

Using BCI in highly immersive VR is a new type of medium. In the experiments described here we have gained a few more insights about the human-interface design factors, which play a critical part in this type of experience.

Our main conclusion from the first experiment is that free choice (or the lack of it) is an important factor, and it affects not only the subjective experience but also task performance. We found out that participants performed better when instructed "what to think," as compared to being free to decide for themselves. This is despite the fact that the free-choice task was very simple and most likely did not involve any significant cognitive effort (none was reported by the participants). It is possible that with additional training, participants could perform BCI tasks with free choice with the same level of accuracy as in the case of the instructed BCI. However, we still consider our experiment as evidence that BCI should be studied in the context of real-usage scenarios.

We believe the method used in this article allows one to compare BCI accuracy indirectly, using task performance, even in the context of free choice. To achieve a more precise comparison between different conditions, tasks, and interaction paradigms, some notion of a baseline BCI performance is required; for example, participants' performance in the training phases. Based on such a baseline, the results of different participants may be normalized. Investigating this normalization requires much more data from more participants.

In the second experiment we have devised a system that allows people to be able to control a virtual body in an IVE, with accuracy ranging from 72 to 96%. We consider this to be a proof of the feasibility of this innovative interface. We have used qualitative methods to get a sense of this new type of experience: What did it feel like? What was the nature of the relationship between the participants and their avatars?

There is growing interest in the BCI community to use IVE, and some evidence that IVE may assist in BCI training, or even improve BCI performance (Leeb et al., 2006; Pfurtshceller et al., 2006). Our finding suggests that BCI in IVE is more enjoyable than traditional BCI, and participants find it more intuitive and natural. However, participants did not seem to perform better when the mapping between their imagery and the feedback was natural, as compared to when this mapping was reversed. This is despite the fact that the participants did report that the reverse condition seemed more confusing and less intuitive. The results we describe in this article thus indicate that the story is complex and justify further research.

In this article we are pointing the way toward the ultimate human-computer interface, an interface through "thought" of a virtual world—as has been described in novels by authors such as William Gibson (1984) and Neal Stephenson (1991). The research described in this article has shown that it is possible to control a virtual body by "thought" and has explored performance-related results and the subjective experience that this entails.

NOTES

Support. This work has been supported by the European Union FET projects PRESENCIA, IST-2001-37927, and PRESENCCIA, IST-2006-27731.

Acknowledgements. We would like to thank Larisa Dikovsky for help in carrying out the second experiment. We thank Marco Gillies and Vinoba Vinoyagamoorthi for allowing us to use the virtual street environment. We are also grateful to David Swapp, Anthony Steed, and other members of the VEGG lab in UCL for their support.

Authors' Present Addresses. Doron Friedman, Advanced Virtuality Lab, Sammy Ofer School of Communications. The Interdisciplinary Center (IDC), P. O. Box 167, Herzliya, 46150, Israel. Email: doronf@idc.ac.il. Robert Leeb, Laboratory of Brain-Computer Interfaces, Institute of Knowledge Discovery, Graz University of Technology, Krenngasse 37, A-8010 Graz, Austria. Email: robert.leeb@epfl.ch. Gert Pfurtscheller, Laboratory of Brain-Krenngasse 37, A-8010 Graz, Austria. Email: pfurtscheller@tugraz.at. Mel Slater, ICREA-Universitat de Barcelona, Facultat de Psicologia, Campus de Mundet- Edifici Teatre, Passeig de la Vall d'Hebron 171, 08035 Barcelona, Spain. Email: melslater@gmail.com

HCI Editorial Record. First manuscript received January 2, 2007. Revision received July 18, 2008. Accepted by Shumin Zhai. Final manuscript received May 3, 2009.—*Editor*

REFERENCES

Armel, K. C., & Ramachandran, V. S. (2003). Projecting sensations to external objects: Evidence from skin conductance response. *Proceedings of the Royal Society of London, Series B, Biological Sciences, 270,* 1499–1506.

Bayliss, J. D. (2003). Use of the evoked potential P3 component for control in a virtual apartment. *IEEE Transactions on Rehabilitation Engineering, 11,* 113–116.

Bayliss, J. D., & Ballard, D. H. (2000). A virtual reality testbed for Brain Computer Interface research. *IEEE Transactions on Rehabilitation Engineering, 8,* 188–190.

Botvinick, M., & Cohen, J. (1998). Rubber hands "feel" touch that eyes see. *Nature, 391*(6669), 756.

Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. (1992, June). The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM, 35*(6), 65–72.

Desmond, D., Horgan, O., & MacLachlan, M. (2002). *Trinity assessment of body plasticity (TABP)*. Unpublished manuscript, Department of Psychology, Trinity College, Dublin, Ireland.

Dornhege, G., Millan, J. R., Hinterberger, T., McFarland, D. J., & Muller, K. R. (Eds.). (2007). *Toward brain-computer interfacing*. Cambridge, MA: MIT Press.

Dragoi, V., & Staddon, J. E. (1999). The dynamics of operant conditioning. *Psychological Review, 106,* 20–61.

Frecon, E., Smith, G., Steed, A., Stenius, M., & Stahl, O. (2001, February). An overview of the COVEN platform. *Presence: Teleoperators and Virtual Environments, 10*(1), 109–127.

Friedman, D., Brogni, A., Guger, C., Antley, A., Steed, A., & Slater, M. (2006). Sharing and analyzing data from presence experiments. *Presence: Teleoperators and Virtual Environments, 15,* 555–610.

Friedman, D., Leeb, R., Guger, C., Steed, A., Pfertscheller, G., & Slater, M. (2007). Navigating virtual reality by thought: What is it like? *Presence: Teleoperators and Virtual Environments, 16,* 100–110.

Gibson, W. (1984). *Neuromancer*. London: Voyager.

Guger, C., Edlinger, G., Harkam, W., Niedermayer, I., & Pfurtscheller, G. (2003). How many people are able to operate an EEG-based brain-computer interface (BCI)? *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 11,* 145–147.

Hettinger, L. J., & Riccio, G. E. (1992). Visually induced motion sickness in virtual environments. *Presence, 1*(3), 306–310.

Jeannerod, M., & Frak, V. (1999). Mental imaging of motor activity in humans. *Current Opinion in Neurobiology, 9*(6), 735–739.

Kelly, F. D., & Main, F. O. (1978). Ideographic research in individual psychology: Problems and solutions. *The Journal of Individual Psychology, 34*(2), 221–231.

Kubler, A., Kotchoubey, B., Kaiser, J., & Wolpaw, J. R. (2001). Brain-computer communication: Unlocking the locked in. *Psychological Bulletin, 127,* 358–375.

Kubler, A., & Muller, K. R. (2007). An introduction to brain-computer interfacing. In *Toward brain-computer interfacing* (pp. 1–26). Cambridge, MA: MIT Press.

Leeb, R., Friedman, D., Müller-Putz, G., Scherer, R., Slater, M., & Pfurtscheller, G. (2007). Self-paced (asynchronous) BCI control of a wheelchair in virtual environments: A case study with a tetraplegic. *Computational Intelligence and Neuroscience, 2007*. doi:10.1155/2007/ 79642

Leeb, R., Keinrath, C., Friedman, D., Guger, C., Neuper, C., Garau, M., et al. (2006). Walking from thoughts: Not the muscles are crucial, but the brain waves! *Presence: Teleoperators and Virtual Environments, 15,* 500–514.

Leeb, R., Lee, F., Keinrath, C., Scherer, R., Bischof, H., & Pfurtscheller, G. (2007). Brain-computer communication: Motivation, aim, and impact of exploring a virtual apartment. *IEEE Transactions on Rehabilitation Engineering, 15,* 473–482.

Lotte, F., Congedo, M., Lecuyer, A., Lamarche, F., & Arnaldi, B. (2007). A review of classification algorithms for EEG-based brain-computer interfaces. *Journal of Neural Engineering, 4,* R1–R13.

Lotze, M., Montoya, P., Erb, M., Hülsmann, E., Flor, H., Klose, U., et al. (1999). Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study. *Journal of Cognitive Neuroscience, 11,* 491–501.

Mackinlay, J. D., Card, S. K., & Robertson, G. G. (1990). Rapid controlled movement through a virtual 3d workspace. *Proceedings of the 17th annual conference on Computer Graphics and Interactive Techniques*. New York: ACM.

Mason, S. G., & Birch, G. E. (2000). A brain-controlled switch for asynchronous control applications. *IEEE Transactions on Biomedical Engineering, 47,* 1297–1307.

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,* 211–214.

Millan, J. D. R., & Mourino, J. (2003). Asynchronous BCI and local neural classi_ers: An overview of the adaptive brain interface project. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 11,* 159–161.

Nelson, W., Hettinger, L., Cunningham, J., & Roe, M. (1997). Navigating through virtual ight environments using brain- body-actuated control. *Proceedings of the IEEE Virtual Reality Annual International Symposium.* IEEE Computer Society.

Neuper, C., Scherer, R., Reiner, M., & Pfurtscheller, G. (2005). Imagery of motor actions: Differential effects of kinesthetic and visual-motor mode of imagery in single-trial EEG. *Cognitive Brain Research, 25,* 668–677.

Pandzic, I. S., Thalmann, N. M., Capin, T. K., & Thalmann, D. (1997). Virtual life network: A body-centered networked virtual environment. *Presence–Teleoperators and Virtual Environments, 6,* 676–686.

Pfurtshceller, G., Leeb, R., Keinrath, C., Friedman, D., Neuper, C., Guger, C., et al. (2006). Walking from thought. *Brain Research, 1071,* 145–152.

Pfurtscheller, G., & Neuper, C. (2001, July). Motor imagery and direct brain computer communication. *Proceedings of the IEEE, 89,* 1123–1134.

Riva, G., David, F., & IJsselsteijn, W. (Eds.). (2003). *Being there: Concepts, effects and measurement of user presence in synthetic environments.* Amsterdam: IOS Press.

Schlogl, A., Kronegg, J., Huggins, J. E., & Mason, S. G. (2007). Evaluation criteria for BCI research. In *Toward brain-computer interfacing* (pp. 327–342). Cambridge, MA: MIT Press.

Schneiderman, B. (1983). Direct manipulation: A step beyond programming languages. *IEEE Computer, 16*(8), 57–69.

Slater, M., Steed, A., McCarthy, J., & Maringelli, F. (1998). The influence of body movement on subjective presence in virtual environments. *Human Factors, 40,* 469–477.

Slater, M., Usoh, M., & Steed, A. (1994, June). Depth of presence in virtual environments. *Presence: Teleoperators and Virtual Environments, 11*(3), 130–144.

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, 2,* 201–219.

Smith, J. (1977). Semi-structured interviewing and qualitative analysis. In J. Smith, R. Harre, & L. Langenhove (Eds.), *Rethinking methods in psychology* (pp. 9–26). London: Sage.

Steed, A., Mortensen, J., & Frecon, E. (2001). Spelunking: Experiences using the DIVE system on CAVE-like platforms. In *Immersive projection technologies and virtual environments* (Vol. 2, pp. 153–164). New York: Springer-Verlag/Wien.

Stephenson, N. (1991). *Snow crash.* New York: Roc.

Sutherland, I. (1965, May). The ultimate display. In *Proceedings of the IFIPS Congress 1965* (Vol. 2, pp. 506–508). Amsterdam, The Netherlands: North-Holland Publishing Co.

Usoh, M., Arthur, K., Whitton, M., Bastos, R., Steed, A., Slater, M., et al. (1999). Walking > walking-in-place > ying, in virtual environments. *Proceedings of SIGGRAPH 99.* New York, NY: ACM Press.

Utz, S. W. (1994). The e_ect of instructions on cognitive strategies and performance in biofeedback. *Journal of Behavioral Medicine, 17,* 291–308.

Vallabhaneni, A., Wang, T., & He, B. (2005). Brain computer interface. In B. He (Ed.), *Neural Engineering* (pp. 85–122). Dordrecht, the Netherlands: Kluwer.

Vidal, J. J. (1973). Toward direct brain-computer communication. *Annual Review of Biophysics and Bioengineering, 2,* 157–180.

Vives, M. S., & Slater, M. (2005). From presence to consciousness through virtual reality. *Nature Reviews Neuroscience, 6*(4), 332–339.

Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., & Vaughan, T. M. (2002). Brain Computer Interfaces for communication and control. *Clinical Neurophysiology, 113,* 767–791.