Is the Rubber Hand Illusion Induced by Immersive Virtual Reality?

Ye Yuan, Anthony Steed[*]

Department of Computer Science, University College London

ABSTRACT

The rubber hand illusion is a simple illusion where participants can be induced to report and behave as if a rubber hand is part of their body. The induction is usually done by an experimenter tapping both a rubber hand prop and the participant's real hand: the touch and visual feedback of the taps must be synchronous and aligned to some extent. The illusion is usually tested by several means including a physical threat to the rubber hand. The response to the threat can be measured by galvanic skin response (GSR): those that have the illusion showed a marked rise in GSR. Based on our own and reported experiences with immersive virtual reality (IVR), we ask whether a similar illusion is induced naturally within IVR? Does the participant report and behave as if the virtual arm is part of their body? We show that participants in a HMD-based IVR who see a virtual body can experience similar responses to threats as those in comparable rubber hand illusion experiments. We show that these responses can be negated by replacing the virtual body with an abstract cursor representing the hand, and that the responses are stable under some gradual forced distortion of tracker space so that proprioceptive and visual information are not matched.

KEYWORDS: Rubber-hand illusion, immersive virtual reality, virtual body, galvanic skin response, body image, body schema.

INDEX TERMS: I.3.7 [Three-Dimensional Graphics and Realism]: Virtual Reality

1 INTRODUCTION

Participants in immersive virtual reality (IVR) systems often behave as if the virtual environments they are experiencing are real [16]. This phenomenon is sometimes referred to as sense of presence [3], but this term has come to mean many different related phenomena in different media. However one class of phenomena are particular to IVR: treatment of virtual stimuli as interacting with one's own body. Participants who are immersed, in the sense that the displays surround and include them, can see a virtual body that closely mirrors the position of their own body. Thus as they move their own body, the visual information they receive from the virtual reality systems, closely mirrors what their proprioception is telling them. The ability to use in-built motor knowledge to access the virtual environment is a key aspect to IVR systems [16].

IVR systems thus afford a unique class of interaction techniques based on knowledge of one's own body size, arm reach, etc. [14]. Indeed, interaction that uses the full body can be easy to learn and effective [22]. Another notable effect is the powerful response users have to virtual drops that are scaled to life size and represented below the feet of the virtual body [13].

In this paper we will make a link between the response to the IVR experience of an interactive virtual body with the illusion known as the rubber hand illusion [2]. In the rubber hand illusion a fake limb can be made to feel as if it is part of your own body. This illusion, with its origins in neuroscience, is used to demonstrate that the brain can be "fooled" by certain types of stimuli, and that one's body image is actually quite malleable.

In an experimental setting, the rubber hand illusion is usually induced by simultaneously tapping the participant's real limb and the fake limb. The taps need to be done simultaneously. The participant is passive during the tapping. The illusion takes a few minutes to occur.

For this paper, the interesting part of the rubber hand illusion is how one tests whether the participant believes that the fake limb is part of his or her body. Typical measures include questionnaires and the stress response of the participant to a perceived threat to the fake limb. We will draw a parallel between these measures and prior work in IVR about the successful use of a virtual body.

Our claim is that an illusion very similar to the rubber hand illusion is "automatically" induced by active use of the virtual body in an IVR. This is a strong claim but we can start to argue for it by using the same evaluation methods that are used to evaluate the rubber hand illusion. Specifically, we can ask the participants about their association with the virtual body, and also threaten their virtual body and look for stress responses.

This paper thus describes an experiment whose protocol is similar to an established rubber hand illusion protocol, but where the passive induction is replaced by an active interaction with the virtual environment. We vary the representation of participant under the hypothesis that if the virtual body follows the participant's action the *IVR arm ownership illusion* will be stronger, than if the participant's body is more abstractly represented. We show that participants that have a realistic virtual body have a higher association with their body than participants that have an abstract representation. This is shown through questionnaires and also their response to a virtual threat which is measured using galvanic skin response (GSR). Further, the magnitude of the response is similar to those demonstrated for the rubber hand illusion.

Further, because of prior work that shows that uniform distortion of proprioception does not hinder effective interaction with IVRs, we show that the illusion does not diminish under slow, uniform distortion of the mapping from tracking coordinates to the virtual hand. Thus, although at the time of threat the virtual hand is displaced 10cm from the real hand, the magnitudes of the responses are not diminished.

The evidence from this experiment by no means proves that the IVR arm ownership illusion is the same as the rubber hand illusion, but we feel that this paper highlights interesting parallels between recent work in neuroscience on body image, and some phenomena that are taken for granted in IVR research.

---

[*] Corresponding author: A.Steed@cs.ucl.ac.uk

2 BACKGROUND

2.1 Rubber Hand Illusion

The rubber hand illusion was first demonstrated by Botvinick and Cohen [2]. A typical induction involves the participant placing his or her hand out of view below a table. On the table a rubber hand is placed. The experimenter simultaneously taps the rubber hand and the real hand synchronously. The participant thus sees the fake hand being touched, but feels their real hand being touched. After a few minutes the participant may start to believe that the fake arm is actually their real arm. In [2] this was assessed by questionnaire, and by having the participant use their other arm (the one not being tapped) to place a mark where they thought their tapped arm was. The magnitude of the displacement from real arm position to mark is taken as a measure of proprioceptive drift. There were strong differences in both the questionnaire and the proprioceptive drift between an experimental condition and a control condition where the taps are asynchronous.

Armel and Ramachandran added another metric to demonstrate the rubber hand illusion: the stress response to a threat to the hand as measured by skin conductance response (galvanic skin response, GSR) [1]. After an induction of 3 minutes, the experimenter lifted a real finger and one rubber finger, but only the rubber one was bent backwards to a position that would appear to be painful. They found that a clear GSR response was generated during this bending.

The illusion is now well-studied in the neuroscience literature. In particular, potential neural mechanisms responsible for integrating the tactile and visual information have been identified [5][7][8]. For the purposes of this paper, one important aspect of the literature is the extent to which the fake arm must look like a real arm. Armel and Ramachandran suggested that the illusion can be generated without an arm: the table-top can simply tapped [1]. Tsakiris and Haggard contradicted this, and suggested that there must be some correlation between the fake arm and the real arm [21]. Thus, although the literature suggests that the fake arm must look like an arm, it does not appear to be important that it look like the participant's own arm with the correct skin colour and garments. This suggests that the illusion may work with a virtual arm.

Ijsselsteijn et al. [11] produced a virtual arm illusion using a projection of an arm on the table. Although this was called a virtual reality induction, the image would have appeared flat to the participant. Slater et al. [20] created a rubber hand illusion in a constrained situation using stereo head-tracked imagery. A participant stood in front of a large display wall, with their right arm held crooked out at shoulder height and hidden behind a screen. A virtual arm was constructed to appear to be pointing straight out from their shoulder. The head was tracked so that the virtual arm would appear solid. The tapping was done using a virtual object touching the virtual arm: this was controlled by a tracker attached to an object touching the real arm. Slater et al. claimed that the magnitude of the response in their experiment was higher than that shown by Ijsselsteijn et al.

More recently, it has been shown using a head-mounted display showing stereo real-time video imagery that an illusion similar to the rubber hand illusion can be induced for the whole body [15].

2.2 Proprioception and Immersive Virtual Reality

Perhaps the key defining feature of IVR is that the systems immerse the participant in the displays [3]. There are two main classes of IVR display technology: physically surrounding displays (e.g. CAVE™-like [5]) and head-mounted displays (HMDs). Only the latter obscure the participant's real body and thus allow a virtual body to be substituted. It has long been argued that a virtual body is a critical component in HMD-based IVRs, and that it has a profound effect on the participant [10][19].

Typically in a HMD-IVR, the virtual body would be closely aligned to the real body. However, Groen and Werkhoven [9] showed that a 10cm offset between tracked position and real position did not hinder participants on a visuo-motor control task. Burns et al. took this further, and introduced more radical distortions between the real and virtual arm positions, again with negligible impact on the ability of the participant to perform visuo-motor tasks [4].

It has been shown that the amount by which the participants use their virtual body can impact their presence responses. Slater and Steed [18] argued that participant who had to use their virtual body to touch objects to activate them had a higher sense of presence than those who simply pressed a button. Slater et al. [17] and later Usoh et al. [22] argued that mimicking walking in an IVR, which creates match between vision and proprioception, leads participants to report a higher sense of presence.

3 EXPERIMENT DESIGN

Based on the review of both the rubber hand illusion literature and the work on IVR, we derived two main hypotheses for this study. First, if the participant actively experiences a visual and proprioceptive match during an IVR experience, they will associate the virtual arm with their own body (*IVR arm ownership illusion*). This association will be tested using a questionnaire and the galvanic skin response (GSR) to a threat. We expect this association will not happen with a virtual body based on a simple abstract arrow cursor that represents the hand's position. Second, we expect that this association will remain even under tracker distortion of a limited amount. That is, there will be no difference between the response of a virtual body where the proprioception and visual information do not exactly match compared to one where they are closely matched. This leads to a two by two design, where the four conditions are: virtual body no distortion, virtual body with distortion, arrow no distortion, arrow with distortion. We use a between subjects design.

3.1 Overview

A standard rubber hand illusion protocol involves a passive induction followed by a short testing phase. We specifically wanted the participant to be active throughout the experiment, and make lots of arm motion to exercise the visuo-proprioceptive matching. This thus constrained us a little in the design of measures for the IVR arm ownership illusion. Specifically the proprioceptive drift test from the standard rubber hand illusion is difficult for us to measure because the active arm is constantly moving. Thus we dropped this measure and relied on the questionnaire and GSR responses (see later). We also needed to include the tracker drift phase, and a threat to the virtual arm.

An overview of the experiment procedure is shown in Figure 1. The IVR part of the experiment lasts 16 minutes. The first three minutes is a baseline period. Then the participant performs a "Simon Game" based on a task in the Burns et al. paper [4], see Section 3.4. In two of the conditions, 2 minutes in to the game, the tracker starts to drift, see Section 3.3. It does this over 3 minutes. The drift offset remains static for the remainder of the task. The participant then switches to a ball game, see Section 3.4. During this game their arm is threatened by a falling lamp. Both games involve the participant using their hand in front of them so that they can see the virtual body, see Section 3.3

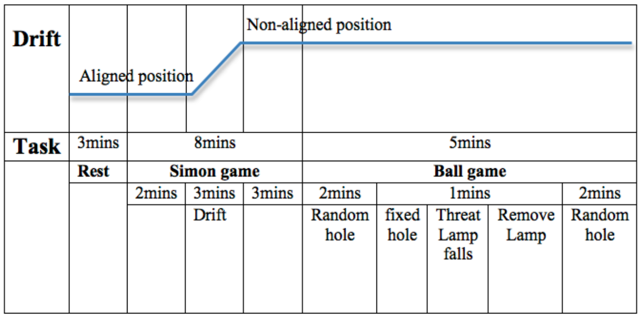

Task	3mins	8mins			5mins				
	Rest	Simon game			Ball game				
		2mins	3mins	3mins	2mins	1mins			2mins
			Drift		Random hole	fixed hole	Threat Lamp falls	Remove Lamp	Random hole

Figure 1: Overview of the experimental protocol, showing (Top) tracker drift for the two conditions with drift and (Bottom) the participant's task.

Figure 2: Participant wearing the VR1280 helmet while sitting in front of a physical table.

3.2 Equipment

The physical configuration of the experiment is shown in Figure 2. The participant was wearing a Virtual Research VR1280 helmet, which has 1280 x 1024 pixel resolution screens and a 60-degree field-of-view with 100% overlap. They were seated in front of a small table.

The graphics were generated by a self-built PC comprising of a dual-core 1.6GHz Intel processors and 2GB main memory. The PC had two graphics cards: one NVidia GeForce 6800 PCI-express card to drive the HMD via two video outputs and one NVidia GeForce 5950 PCI card to drive a control screen.

Tracking information was generated by an Intersense IS-900 system. The participant sat in a CAVE™-like system, the UCL ReaCTor, but this was solely so that the tracking system did not need to be moved. One tracker was placed on the front of the HMD. The participant held a wand tracker in his or her right hand.

GSR data was recorded by a NeXus-4 device. The two sensors were fitted to two fingers of the left hand, which was passive during the experiment. The participant was asked to leave this hand on the table to reduce any artifacts in the GSR from motion. GSR data was recorded using the Biotrace+ software supplied with the NeXus-4. Biotrace+ was running on a second PC.

The experiment was implemented in the XVR software from VRMedia. The XVR software and the Biotrace software produced separate log files. The clocks of the two PCs were synchronized prior to the experiment and checked frequently. Analysis was done in MATLAB® 2008b.

Figure 3: A 3rd person view of the avatar of the participant inside the virtual environment. The virtual table is registered to the real table.

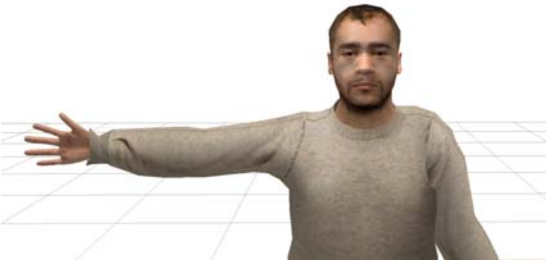

Figure 4: The avatar that the participant sees from a 1st person view. The right arm of the avatar is animated with joints at shoulder, elbow and wrist.

3.3 Virtual Body & Tracker Drift

The virtual environment scenario is shown in Figure 3 from a third person point of view. Participants would see a virtual body from a first person point of view. The avatar is shown in more detail in Figure 4. We integrated an inverse kinematics system for the right arm of this avatar. The avatar has joints at the shoulder, elbow and wrist. The free degree of freedom in the elbow was constrained so that the elbow was as low as possible, but above the virtual table that was registered to the real table. The participant was instructed to hold the wand in a specific manner (see Figure 2). The virtual hand was attached to the wand mimicking the orientation of the real hand. The virtual hand was either holding a remote control or was posed in a grasp gesture. In the two drift conditions, the virtual hand was offset horizontally to the right of the participant. This meant that the participant would have to hold their hand further to their left in order to point at targets or pick objects. The drift was achieved by a simple world-coordinate offset in the coupling of the tracker position to the hand position of the avatar. The virtual hand or arrow was moved with a speed of 0.56 mm/s. After 3-minute drifting period, the displacement between real and virtual hand would be 10cm

Although the visual appearance of the avatar is male, the participant would not see the head, and the avatar's hand was not noticeably masculine or feminine in appearance. In the two conditions that use an arrow the right arm is not shown, and a 20cm long arrow is shown, see Figure 7. The arrow is placed centrally at the current tracker position, which is in the centre of the wand device. The tip of the arrow is used to select objects. The left arm is always depicted in the virtual environment, but is static throughout the experiment. Because the participant wears the GSR sensors on their left arm, they are instructed not to move it during the experiment.

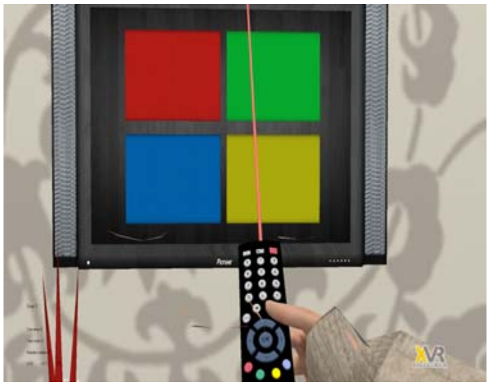

Figure 5: The virtual Simon game. The avatar is seen holding a remote control which they point at coloured squares on a virtual monitor on the wall.

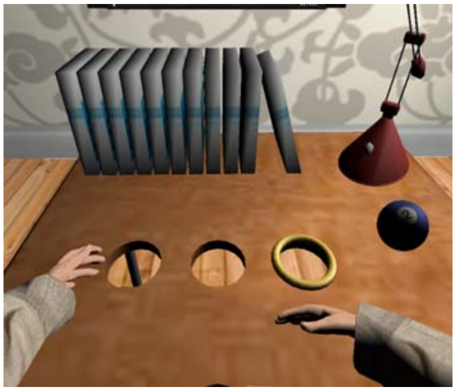

Figure 6: The ball dropping task. The ball appears on the right and the participant must pick it up and drop it through the hole that has a ring around it.

3.4 Game Tasks & Threat

The participants played two games during the experiment. The first was a Simon-like game, based on memorizing sequences of flashes of color panels. This game was based closely on a game implemented in Burns et al.[4]. The second game was a ball game.

In the Simon game, a virtual screen, see Figure 5, was placed on the wall in front of the participant, see Figure 3. One of the panel on the screen flashes for 1s, and the participant must point at this screen and press a button on the wand to indicate the panel. Two panels then flash for 1s each, 1s apart, and the participant must indicate both in order. The length of the sequence of flashes carries on increasing but only up to five, when it then resets to one making the game easy enough that no-one should make any mistakes. In our experiment no-one did. In our implementation a ray emerges from a virtual remote-control device that the participant appears to hold making it easier to aim. It is during this game that in two of the conditions the virtual will drift away from the tracker position.

In the ball game, the participant must pick up a ball that appears somewhere on the table well within arm's reach and then drop it in one of the three holes on the table, see Figure 6 and Figure 7. Once the ball drops through the hole, it is immediately replaced on the table. The randomly chosen target hole is highlighted by a ring.

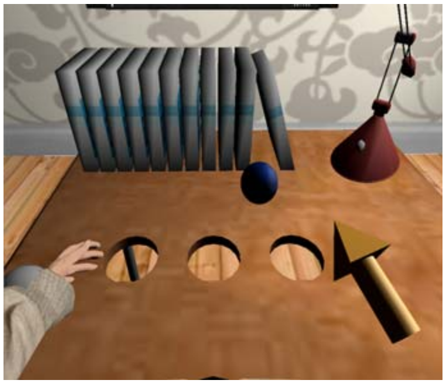

Figure 7: The ball dropping task showing the arrow for the hand.

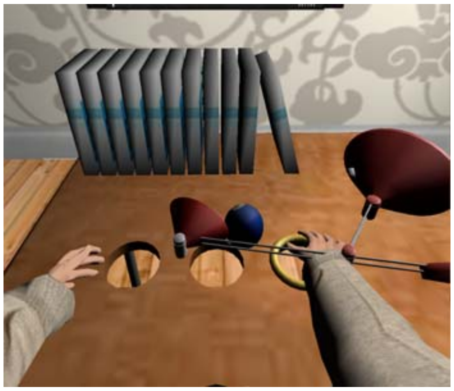

Figure 8: The lamp falling over threatening the virtual hand.

Before the threat is made, a specific hole is highlighted and a specific ball position is chosen. These are both on the right of the table. A table lamp then falls hitting the target hole, the ball, and, hopefully, the participant's virtual hand. The lamp did hit the participant's hand in all our trials, but note that the exact hand position was not constrained.

Finally, 20s after falling, the lamp vanishes. The ball game continues using random targets and the participant carries on until the 16 minutes is up.

3.5 Measures

3.5.1 Questionnaire

We used a variant of the 9-question survey introduced in [2] and modified in [20] to apply to virtual scenarios. The order of questions was changed, and we added one question which we do not present as it gave no insights. Thus below and in the results there is no Question 2 (Q2). Note that the wording of our questions is not the same as [2] or [20] because the tasks and the experience are not the same. Thus, it is not possible to compare directly the results with previous studies.

Participants indicated the strength of agreement for each the 9 statements on a 7-point Likert scale ranging from 1 (=strongly disagree) and 4 (= neutral) to 7 (=strongly agree). There are three statements that were designed to correspond to the illusion: Question 3, 4 and 7:

3. During the experiment there were moments in which I felt as if the virtual arm/arrow was my own arm.

4. Sometimes I had the feeling that I was holding the virtual object (Balls or TV control) in the location of the real arm.

7. During the experiment there were moments in which it seemed that my own arm was being hit by the falling lamp.

Note that the wording of question 7 is necessarily different than the wording of the corresponding question in a rubber hand illusion question because it usually concerns the passive induction. For example in [20] the wording is "*Sometimes I had the feeling that I was receiving the hits in the location of the virtual arm*.). However, we consider 7 to be the corresponding indication of the IVR arm ownership illusion.

The remaining statements are designed as control statements, which are unrelated to the illusion: Question 5,6,8,9 and 10. We dropped one control question from [2] and [20] because it is related to the induction, and we don't have an explicit induction phase. The five control questions are thus:

5. During the experiment there were moments in which it seemed that the virtual object I held was in some place in between my own hand and the virtual hand.

6. During the experiment there were moments in which I felt as if the real hand/arrow was becoming virtual.

8. During the experiment there were moments in which the virtual arm/arrow started to look like my own arm in some aspects.

9. During the experiment there were moments in which I had the sensation of having more than one right arm.

10. During the experiment there were moments in which it seemed that my real arm was being displaced towards the left (towards the virtual arm/arrow).

We added a single question about immersion and presence:

1.During the experiment, how immersed did you feel being in the virtual reality

3.5.2 GSR

The period of interest for the GSR is immediately after the threat from the falling lamp. Our analysis closely follows that in [1], and we focus on the GSR rise in the 5 seconds following the threat.

3.6 Participants & Protocol

Twenty healthy participants, 7 male, 13 female, between 20 and 26 years of age were recruited for the experiment by advertisement posters and emails. Most participants had higher education background and were studying in different universities (6 Computer Science master students, the rest from Architecture, Business Management, and Economy departments). They were each offered £5 for their participation. The information sheet, which described the details of the experiment were sent to each via email after they confirmed their attendance and they were asked to read it before arriving at the laboratory. Participants were randomly assigned to the four conditions, with 5 participants in each condition. This study was approved by the University College London Ethics Committee.

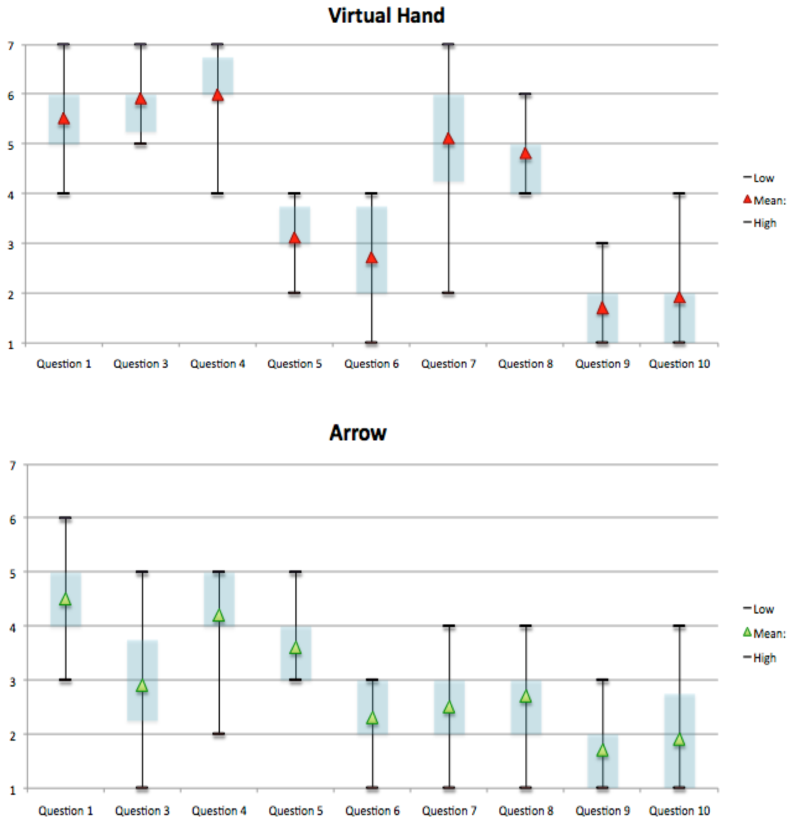

Figure 9: Boxplots for questionnaire results comparing both Virtual Hand conditions (10 subjects) and both Arrow conditions (10 subjects). Triangles indicate mean rating, Boxes indicate the inter-quartile ranges and Bars indicate rating range. Q3,4,7 are the illusion statements. Q5,6,8,9,10 are the control statements.

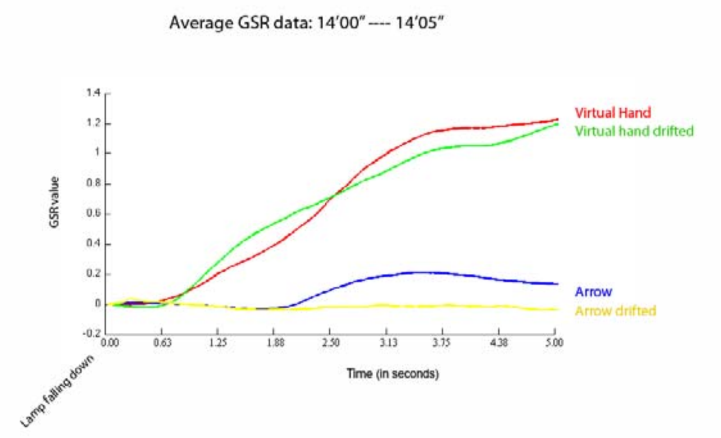

Figure 10: Average GSR responses (microsiemens) for each of the four conditions for the period 0-5 seconds from the threat at 14 minutes.

On arrival the participants were asked to sign a consent form. The equipment and tasks were explained and the participant introduced to the system. They were equipped with the GSR sensor and HMD, and allowed to practice with the wand and asked to verify that they noticed the correspondence between motion and visual response. Then the 16 minute experiment task was started. Immediately after completing the task, participants completed the 10 question questionnaire. This was followed by a short informal interview and debriefing. The whole process took 30-40 minutes.

4 RESULTS

In general there is no significant impact of the distortion of tracking, so in the main we present the results comparing only two Virtual Hand conditions versus the two Arrow Conditions.

4.1 Questionnaire

There was no significant difference between the four conditions on Q1 which concerns immersion in the IVR, though the mean is higher for Virtual Hand (see Figure 9).

For the two Virtual Hand conditions but not the two Arrow conditions, the difference in ratings between the illusion statements (Question 3,4,7) and the control statements (Question 5,6,8,9,10) was significant (two tailed t-test, t-value = 3.2012, df = 19, p = 0.0028; after correction for multiple comparisons). Note also that the ratings for all subjects are low for the control statements (Figure 9).

Q3 is the best discriminator between the Virtual Hand and Arrow conditions.

There were no differences between the ratings of the two Virtual Hand conditions nor the two Arrow conditions. In particular, note that Q10 and to some extent Q5, asks about displacement of objects or limbs, but none of the participants rated these answers highly.

4.2 GSR

The GSR responses of the participants are summarized across conditions in Figure 10. Following [1], a summative measure is taken by finding the maximum rise (max_rise) in the amplitude of the GSR for each subject over the 5 second period after the threat, and then taking log(max_rise + 1)). This is shown in Table 1.

Table 1. Mean and standard error of the mean (SEM) GSRs (log(max_rise + 1)) for all 4 conditions.

Condition	mean (SEM)
Virtual Hand	0.342 (0.06)
Virtual Hand Drifted	0.328 (0.08)
Arrow	0.091 (0.07)
Arrow Drifted	0.083 (0.08)

A comparison of drift and no drift Virtual Hand conditions on GSR shows that they were not significantly different (t = 0.4965, df = 4, p = 0.637). Similarly for the drift and no drift Arrow (t = 2.2998, df = 4, p = 0.562).

Comparing both Virtual Hand conditions against both Arrow conditions does show a significant difference (t = -2.8505, df = 9, p = 0.011).

4.3 Debriefing and Observations

Debriefings were held with all participants. One of the most interesting comments was one participant that reported that illusion was convincing that he found himself wondering why he wore long-sleeve sweater in summer (this was what the avatar's was shown wearing). Furthermore, during the experiment, we observed that two participants pulled their real hand away to dodge the falling lamp. In the debriefing one out of the 20 subjects reported feeling pain when the virtual hand was threatened.

5 DISCUSSION

The results confirmed the initial hypotheses of the study. By the use of a questionnaire and the GSR response to a threat, we found a significant difference between the Virtual Arm and Arrow conditions, with the Virtual Arm showing a strong response to the threat and questionnaire on perception of arm ownership, and the Arrow not. We also found that the small distortion of tracking registration did not impact the response.

The results of the questionnaire concur with similar findings in studies on the original rubber hand illusion effect. Furthermore, we have shown that the IVR arm ownership illusion appears to exist when the virtual arm roughly appears in shape and animation like the participant's own arm, but not when there is a virtual arrow.

We have not performed the usual test of proprioceptive drift that is done in rubber hand illusion tests. Partly this is because in our two non-drift conditions, there should be zero proprioceptive drift: the virtual arm and the real arm are actually in the same place. Note that in the rubber hand illusion, the presence of the illusion is marked by this measure by tendency to indicate that their right arm is where the fake arm is, not their real arm. Of course, we have actually created an offset in two of our conditions, and it would be interesting to ask whether the participant's still understood that their real arm was not in the same position as the virtual arm. We did not include a method for assessing this, but note that our participants were, as in previous experiments on tracker distortion, very able to interact with the virtual environment successfully. Although we did not include it explicitly as a measure, log files from the experiment show no difference in the rate at which the participants completed the Simon game or ball game depending on tracker distortion or not. This doesn't mean that there isn't an effect, but our experiment was not designed to elucidate this.

The GSR response is quite significant in the two virtual hand conditions. The results can be compared with [1] where the GSR test was introduced. In that, the maximum mean SCR response for the basic rubber hand illusion is 0.39 (0.07 SEM). They achieve 0.45 (0.06 SEM) in another protocol, and in a control condition of 0.11 (0.04 SEM). Thus, we could hypothesize that the IVR arm ownership illusion is not as strong. The strength of the response might be increased by better representation or interaction, or allowing the participant more freedom to move. Finally, the virtual arrow appears to be a good control condition as it allows the participant to interact successfully, but doesn't appear to lead to an ownership illusion similar to that experience with a virtual hand.

6 CONCLUSIONS & FUTURE WORK

In this paper we have shown that an "IVR arm ownership illusion" exists and that it can be tested for and measured using a protocol derived directly from those described in the literature on the rubber hand illusion. The evidence from our experiment by no means proves that the IVR arm ownership illusion is the same as the rubber hand illusion, but we feel that this paper highlights interesting parallels between the neuroscience-based work and the history of the study of effective IVR.

The results lend more weight to the argument that a virtual body is an important component of an IVR experience. As shown under different circumstances and different setups (e.g. [12][19]), the presence of a virtual body with arms can significantly alter how the participant interacts with the virtual world. Our results thus suggest that the IVR arm ownership illusion might be a good test of the effectiveness of a virtual body or might be used as a proxy for immersion or engagement in an IVR experience.

In a HMD-based IVR it would be interesting to investigate how the representation and behavior of the virtual body affects the response. For example, one could track the arm more exactly using motion capture, or one can imagine building systems that more significantly distort the mapping of tracking space to virtual body representation.

Further, it would be interesting to see if a similar illusion can be elicited in non-HMD IVRs: of course, it isn't possible in a CAVE-like display to have something fall on the arm, but other threats might be possible.

ACKNOWLEDGEMENTS

We would like to acknowledge Bernhard Spanlang of Universitat de Barcelona for supplying the avatar used in the experiment and Angus Antley of UCL for providing inverse kinematics scripts for XVR and for helping us to integrate them in to our experiment. We would like to acknowledge the support of the European Union FET project PRESENCCIA Contract Number 27731. We would also like to thank David Swapp of UCL for supporting the experiment.

REFERENCES

[1] K. C. Armel and V. S. Ramachandran. Projecting sensations to external objects: Evidence from skin conductance response. Proceedings of the Royal Society of London: Biological, 270, 1499–1506, 2003.

[2] M. Botvinick and J. Cohen. Rubber hands 'feel' touch that eyes see. Nature 391, 756, 1998.

[3] K-E. Bystrom, W. Barfield and C. Hendrix. A Conceptual Model of the Sense of Presence in Virtual Environments, Presence: Teleoperators and Virtual Environments, 8(2), 241-244, 1999.

[4] E. Burns, S. Razzaque, A. T. Panter, M. C. Whitton, M. R. McCallus and F. P. Brooks. The hand is slower than the eye: a quantitative exploration of visual dominance over proprioception. In Proceedings of IEEE Virtual Reality 2005, 3-10, 2005.

[5] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon and Hart, J. C. The CAVE: audio visual experience automatic virtual environment. Communications of the ACM 35, 6 (Jun. 1992), 64-72, 1992.

[6] H. Ehrsson, N. P. Holmes and R. E. Passingham. Touching a rubber hand: feeling of body ownership is associated with activity in multisensory brain areas. Journal of Neuroscience, 25, 10564–10573, 2005.

[7] H. Ehrsson, C. Spence and R. E. Passingham. That's my hand! Activity in premotor cortex reflects feeling of ownership of a limb. Science, 305, 875–877, 2004.

[8] M. S. Graziano. Where is my arm? The relative role of vision and proprioception in the neuronal representation of limb position. Proceedings of the National Academy of Sciences, USA, 96, 10418–10421, 1999.

[9] J. Groen and P. J. Werkhoven. Visuomotor adaptation to virtual hand position in interactive virtual environments. Presence: Teleoperators and Virtual Environments, 7, 5, 429-446, 1998.

[10] C. Heeter. Reflections on Real Presence by a Virtual Person, Presence: Teleoperators & Virtual Environments, 12, 4, 335-345, 2003.

[11] W. IJsselsteijn, Y. de Kort and A. Haans. Is this my hand I see before me? The rubber hand illusion in reality, virtual reality and mixed reality. Presence: Teleoperators & Virtual Environments, 15, 4, 455–464, 2006.

[12] B. Lok, S. Naik, M. Whitton and F. P. Brooks. Effects of Handling Real Objects and Avatar Fidelity on Cognitive Task Performance in Virtual Environments. In Proceedings of IEEE Virtual Reality 2003, Los Angeles, CA (March 22-26), 125-132, 2003.

[13] M. Meehan, B. Insko, M. C. Whitton and F. P. Brooks. Physiological Measures of Presence in Stressful Virtual Environments, ACM Transactions on Graphics, 21, 3, 645-652, 2002.

[14] M. R. Mine, F. P. Brooks and C. H. Sequin. Moving objects in space: exploiting proprioception in virtual-environment interaction. In Proceedings of the 24th Annual Conference on Computer Graphics and interactive Techniques International Conference on Computer Graphics and Interactive Techniques. ACM Press/Addison-Wesley Publishing Co., New York, NY, 19-26, 1997.

[15] V. Petkova and H. Ehrsson. If I Were You: Perceptual Illusion of Body Swapping, PLoS ONE, 3, 12, 2008.

[16] M. V. Sanchez-Vives and M. Slater. From presence to consciousness through virtual reality. Nature Reviews Neuroscience, 6, 332–339, 2005.

[17] M. Slater, M. Usoh and A. Steed. Taking Steps: The Influence of a Walking Metaphor on Presence in Virtual Reality, ACM Transactions on Computer- Human Interaction, 2, 3, 201-219, 1995.

[18] M. Slater and A. Steed. A Virtual Presence Counter. Presence: Teleoperators & Virtual Environments, 9, 5, 413-434, 2000.

[19] M. Slater, A. Steed, J. McCarthy and F. Marinelli. The Influence of Body Movement on Presence in Virtual Environments, Human Factors: The Journal of the Human Factors and Ergonomics Society, 40, 3, 469-477, 1998.

[20] M. Slater, D. Perez-Marcos, H. Ehrsson and M. V. Sanchez-Vives. Towards a digital body: the virtual arm illusion. Frontiers in Human Neuroscience, 2, 6, 2008.

[21] M. Tsakiris and P. Haggard. The rubber hand illusion revisited: visuo- tactile integration and self-attribution. Journal of Experimental Psychology: Human Perception and Performance, 31, 80–91, 2005.

[22] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater and F. P. Brooks. Walking > walking-in-place > flying, in virtual environments. In Proceedings of the 26th Annual Conference on Computer Graphics and interactive Techniques International Conference on Computer Graphics and Interactive Techniques. ACM Press, New York, NY, 359-364, 1999.