Is This *My* Hand I See Before Me?
The Rubber Hand Illusion in Reality, Virtual Reality, and Mixed Reality

W. A. IJsselsteijn, Y. A. W. de Kort, A. Haans
Eindhoven University of Technology, Human-Technology Interaction Group
{w.a.ijsselsteijn@tue.nl, y.a.w.d.kort@tm.tue.nl, a.haans@tm.tue.nl}

Abstract

This paper presents a first study in which a recently reported intermodal perceptual illusion known as the rubber hand illusion is experimentally investigated under mediated conditions. When one's own hand is placed out of view and a visible fake hand is repeatedly stroked and tapped in synchrony with the unseen hand, subjects report a strong sense in which the fake hand is experienced as part of their own body. In our experiment, we investigated this illusion under three conditions: (i) unmediated condition, replicating the original paradigm, (ii) virtual reality (VR) condition, where both the fake hand and its stimulation were projected on the table in front of the participant, and (iii) mixed reality (MR) condition, where the fake hand was projected, but its stimulation was unmediated. Dependent measures included self-report (open-ended and questionnaire-based) and drift, that is, the offset between the felt position of the hidden hand and its actual position. As expected, the unmediated condition produced the strongest illusion, as indicated both by self-report and drift towards the rubber hand. The VR condition produced a more convincing subjective illusion than the MR condition, although no difference in drift was found between the mediated conditions. Results are discussed in terms of perceptual mechanisms underlying the rubber hand illusion, and the illusion's relevance to understanding telepresence.

Keywords--- **Rubber hand illusion, multisensory integration, body image, virtual reality, mixed reality**

1. Introduction

One need not be a chamber to be haunted,
One need not be a house;
The brain has corridors surpassing
Material place

– Emily Dickinson

Human brains seem to support highly malleable body images. Although intuitively we expect our body image to be durable and permanent, evidence is mounting that suggests that our sense of bodily self-identification – the ability to distinguish what's contained within versus what's beyond our familiar biological shell – is a flexible, temporary construct and not a fixed property. Having a negotiable body image has clear survival value when considering the profound bodily changes that the brain has to accommodate during a lifetime of physical development and change. What is most surprising here, however, is the relative speed at which the brain appears to support a significantly altered body image after just a few minutes of the right kind of sensory stimulation.

A particularly interesting and relevant phenomenon in this respect is a recently reported intermodal perceptual illusion known as the rubber hand illusion [1,2,3]. When a person is watching a fake, rubber hand being stroked and tapped in precise synchrony with his or her own unseen hand, the person will, within a few minutes of stimulation, start experiencing the rubber hand as an actual part of his or her own body. In part, this illusion illustrates the importance of visual information in specifying limb location and constructing the body image (cf. [4]). For example, when seen and felt hand position are in conflict, as is the case when one wears a prism that displaces the entire visual field to one side, the visually displaced hand is usually felt where it is seen, a phenomenon know as immediate visual capture [5]. The visual adaptation of proprioceptive position that occurs during the rubber hand illusion is related, though not identical, to prism adaptation (see [6] for an overview). After prolonged exposure to prism-induced visual displacements, after-effects will occur including misreaching in the direction opposite to the previous visual displacement. Similar effects have been reported in adapting to tele-systems and virtual environments (see, e.g., [7]). However, a key distinguishing feature of the rubber hand illusion is that it emerges from closely correlated visual and tactile stimulation, resulting in a strong sense of body ownership of the fake hand. The correlation between visual, tactile and proprioceptive information can be thought of as self-specifying for bodily self-identification, as the brain has learned from a very early age onwards that it can only be the body, and no other object, that can register these specific intersensory correlations [8].

The extent to which non-biological artefacts, such as a rubber hand, can be incorporated as a phenomenal extension of the self has clear relevance to the area of telepresence [9]. Understanding the conditions under which such integration may or may not occur has implications for the design of virtual environments, teleoperation and mixed reality systems, and ways in which the body may be optimally represented in such mediated environments. More importantly, it enhances our fundamental understanding of the phenomenal experience of telepresence and the psychological and brain mechanisms involved in distinguishing self from non-self, and reality from mediation.

In this paper, we report on an experiment we performed to investigate the rubber hand illusion under mediated conditions. However, before describing the rationale of the experiment, we will first turn to the rubber hand illusion in more detail.

1.1. The Rubber Hand Illusion

Botvinick and Cohen [1] provided a first description of the rubber hand illusion. This crossmodal perceptual illusion occurred when participants' left hand was placed out of view and a life-size rubber facsimile of a human hand was placed in front of them. Subsequently, both the rubber hand and participants' left hand were gently stroked by two small paintbrushes, synchronizing timing as closely as possible. Subjects reported feeling a sense of ownership of the rubber hand, as if it was actually their own. In addition to self-report, Botvinick and Cohen also employed a measure of drift, where subjects were asked to close their eyes and align their right index finger with the index finger of their unseen left hand. Results showed a proprioceptive drift towards the rubber hand, with the magnitude of drift correlating significantly with the reported duration of the illusion.

Although Botvinick and Cohen interpret their results as an effect of visual information overriding the incongruent proprioceptive information, Armel and Ramachandran [2] contest this claim, demonstrating that the illusory sensation can also be elicited by merely stimulating the tabletop in front of a participant, which bears no visual resemblance to a hand (see also [10]). They argue that the illusion mainly arises "from the 'Bayesian logic' of all perception; the brain's remarkable ability to detect statistical correlations in sensory inputs in constructing useful perceptual representations of the world – including one's body."([2], p. 1500). Armel and Ramachandran (2003) further showed that when the physical integrity of the rubber hand was threatened (bending a finger backwards to seem painful), a clear skin conductance response was generated. The illusion could even be projected to anatomically impossible locations, with the rubber hand positioned at a distance. It is important to note, however, that although Armel and Ramachandran's study showed that the rubber hand illusion is relatively robust to manipulations of form or location (i.e., the illusion still occurs to an extent), the subjective intensity appears to be much lower under these circumstances, and in particular in the tabletop condition. This questions the authors' interpretation that the illusion is resistant to top-down knowledge from cognitive body representations and is solely governed by the brain's ability to extract statistical correlations when perceptions from different modalities co-occur with a high probability. Indeed, a series of experiments recently reported by Tsakiris and Haggard [3] support the contention that bottom-up visuotactile correlations are modulated by top-down influences originating from one's body representation in creating the rubber hand illusion. However, Tsakiris and Haggard's results are solely based on measuring drift, making direct comparisons between their results and those of Armel and Ramachandran difficult.

1.2. Rationale of the Current Experiment

The experiment reported in this paper was performed for three reasons. First of all, we wanted to introduce intermediate levels in form manipulation between the original rubber hand illusion as reported by Botvinick and Cohen [1] and the 'table illusion' as reported by Armel and Ramachandran [2]. Teasing apart and testing the various form factors that influence the vividness of the rubber hand illusion will allow us to better understand the contributing processes, in particular the role of the cognitive body representation, underlying the illusion. To this end, we chose to use a video-projection of a rubber hand (and its synchronous stimulation) onto the flat tabletop surface (we dubbed this the 'virtual reality' condition), thus reproducing the rubber hand form in terms of basic contour, size, texture and colour. The main perceptual difference was in terms of perceived rubber hand volume. By using a non-tracked, monoscopic projection, the stereoscopic and motion parallax cues to object shape were absent, allowing us to assess the impact that these cues have in activating our cognitive body scheme by comparing this condition with the unmediated condition, where these cues are available.

Secondly, Armel and Ramachandran [2] reported anecdotally that the table illusion was more vivid if subjects could see a common texture being synchronously manipulated – in their case a band-aid placed on both the subject's real hand and the table surface. To test this, we chose to project the rubber hand on the tabletop in front of the participant (as before), however with the touch stimulation being unmediated, that is, applied directly to the tabletop projection visible in front of the participant instead of to the rubber hand which was being recorded. Thus, this 'mixed reality' condition would allow us to check whether inconsistencies in perceived texture would diminish the vividness of the illusion.

Lastly, since the rubber hand illusion appears to be a cognitively impenetrable perceptual illusion, the level to which it can be reproduced under mediated conditions may provide us with an interesting indicator of the perceptual quality of a particular form of mediation, and thus a potential indicator of presence.

2. Method

2.1. Design

In this study, we aimed to compare the 'traditional' unmediated rubber hand condition (see Figure 1A) with 2 types of mediated conditions. First, in what we call the Virtual Reality condition (VR), both the rubber hand and the stimulation of the rubber hand (via a small painter's brush held by the experimenter) were projected on the table in front of the participant (see Figure 1B), thus providing a fully mediated equivalent of the original rubber hand experiment, as reported by Botvinick & Cohen [1], and employed in subsequent studies by various others. Second, in the Mixed Reality condition (MR), the rubber hand was again projected in front of the participant (as in the VR condition), yet this time the stimulation by the brush was physically applied to the projection of the rubber hand, rather than to the rubber hand itself (see Figure 1C).

As the existing literature points to significant variations between individuals in both the experienced nature and vividness of the rubber hand illusion, we decided to use a basic within-subjects design to control for this potential variation and increase our experiment's sensitivity. To compensate for potential order effects, the three conditions were presented in fully counterbalanced order.

2.2. Participants

Twenty-four participants, 15 male, 9 female, between 20 and 32 years of age, all with normal or corrected-to-normal vision, took part in this study. Twenty participants were right-handed, three were left-handed, and one had mixed handedness. Participants were either students or employees at the Eindhoven University of Technology in Eindhoven, The Netherlands. They were naïve to the hypothesis under test. Students were compensated with € 7 for their participation.

2.3. Setting and Apparatus

The experiment was conducted at the UseLab of the Human-Technology Interaction Group. The UseLab is a usability laboratory equipped with standard living room furniture as well as state-of-the-art observational technologies and tools. Figure 1 shows the setup that was used in this experiment for the three conditions (A: unmediated, B: virtual reality, and C: mixed reality). The fake "rubber" hand used in all conditions was highly realistic in terms of colour, skin texture, size and shape. It was originally developed by Otto Bock Benelux B.V. as a prosthetic left hand and kindly donated to the authors for research purposes.

A wooden separating screen was used to obscure the view the participants had onto their own left hand. Also, in the VR and MR conditions, the rubber hand was itself placed out of view, behind the separating screen. The rubber hand, or its projection, was placed in a natural position in relation to the participant's torso, slightly left in front of the participant. This would be a comfortable position if it were the participant's own hand (i.e., not an anatomically implausible location – cf. [2]). The distance between the participant's left hand, placed out of view, and the rubber hand (or its projection) was approximately 30 cm. Two small brushes were used to synchronously stroke congruent positions on both the rubber hand (or its projection) and the participant's unseen left hand.

A standard mini-DV camera, mounted on a tripod, was used to record the rubber hand and the stimulation in the VR condition, or only the rubber hand in the MR condition. The camera was mounted such that it had a top view of the recording area, on the left side of the separating screen, as depicted in Figure 1 (panels B & C). The camera's output was connected to an InFocus LP750 projector, which projected directly onto the tabletop surface in front of the participant. Care was taken that the rubber hand projection was of the same size as the rubber hand itself, and that its perspective was matching the participant's viewpoint.

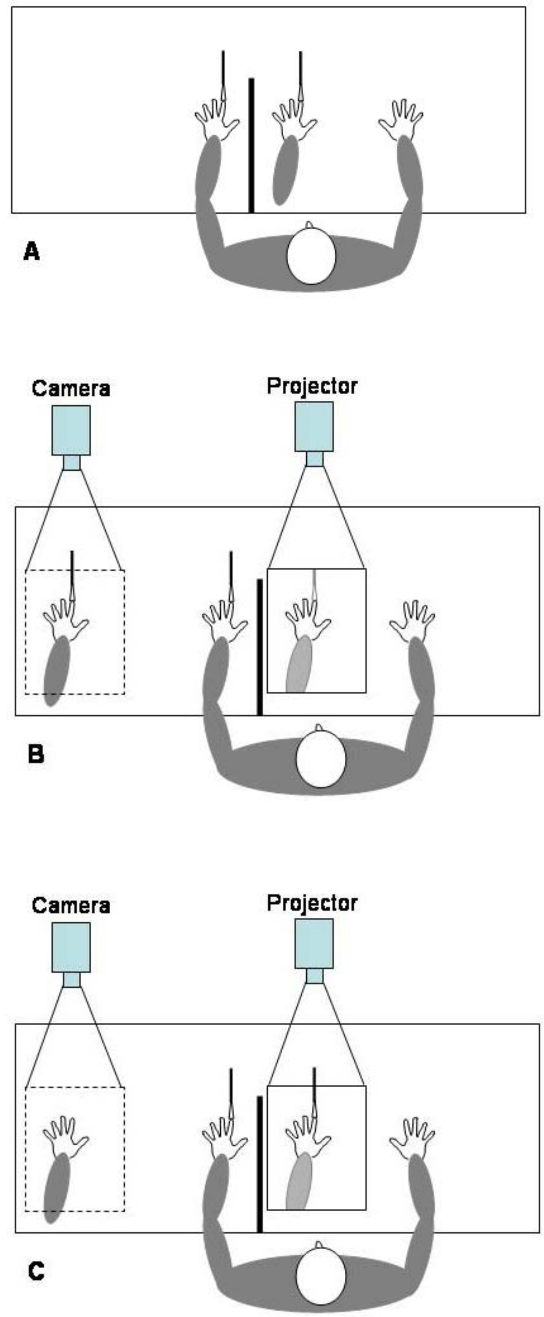

Figure 1. Overview of the three experimental conditions: A) *Unmediated condition*: the rubber hand and its stimulation are both physically present in the participant's field of view. B) *Virtual Reality condition*: both the rubber hand and its stimulation are presented as projections on the table surface in front of the participant. C) *Mixed Reality condition*: the rubber hand is projected in front of the participant and unmediated stimulation is applied to this projection.

2.4. Measurement

In the present experiment we employed self-report to directly assess participants' experiences, and measured drift as an objective corroborative measure of the rubber hand illusion. Self-report included a questionnaire as well as an open-ended, qualitative description of the experience.

2.4.1. Questionnaire The questionnaire was adopted from Botvinick and Cohen [1]. Their questionnaire consists of 9 statements describing specific perceptual effects associated with the rubber hand illusion, such as "I felt the rubber hand was my hand" or "It seemed as though the touch I felt was caused by the paintbrush touching the rubber hand". All items were translated into Dutch.

Three changes were made to this questionnaire. Firstly, the last item ("The rubber hand began to resemble my own (real) hand, in terms of shape, skin tone, freckles or some other visual feature") was divided into two separate items, one tapping resemblance between the rubber hand and the real hand in terms of shape, the other in terms of texture. Secondly, one item was added describing a sensation that a number of people reported during the pilot phase of the study: "It felt as if my hand was inside the rubber hand". Lastly, the 7-point response scale used by Botvinick and Cohen, running from '---' via '0' to '+++' was reformulated to run from 'not at all' to 'completely'. The resulting 11 items are reported in the caption of Figure 2 in the results section.

2.4.2. Drift Drift is a measure gauging a distortion of proprioception in participants that typically occurs after exposure to the rubber hand stimulation. With eyes closed and keeping their left hand in place on the table, participants were asked to indicate the location of their left hand by moving their right hand in a straight line below the table until they feel both hands are in alignment with each other. This task was performed before and after each condition. Drift was calculated by subtracting the displacement to the right (i.e., towards the rubber hand) pre-exposure from the displacement post-exposure (similar to the method used in [3]).

2.5. Procedure

As our study was aimed to elucidate to what extent the rubber hand illusion would occur under mediated conditions, we selected participants on the basis of a short pilot test in which it was established that they indeed were able to experience the rubber hand illusion. Of the 30 participants that were recruited, 6 were excluded from partaking in the study as they did not report any sign of the illusion. The pilot test and the main experimental study were at least one week apart.

On arrival at the UseLab, participants were seated behind a standard office table with a white tabletop surface, and were asked to place their left hand palm face down in a relaxed position on top of a marker behind the wooden partition. This setup ensured that participants were unable to view their real left hand and arm. Participants were instructed not to move their left hand during the experiment, and to focus their attention on the fake hand that was placed in a natural position in front of them.

The experiment was divided into three sessions, one for each condition. Conditions were completely counterbalanced, yielding 6 unique orders. In the unmediated and VR conditions, the experimenter synchronously stroked the fingers of the participant's invisible left hand and the rubber hand for approximately 7,5 minutes, using a small brush. In the MR condition, the experimenter stroked the *projection* of the rubber hand on the table surface in front of the participant instead of the rubber hand itself. After 7,5 minutes of synchronous stimulation in each condition, participants were asked to immediately close their eyes and indicate the felt position of their left hand, in order to establish a measure of drift. Subsequently, participants were asked to fill out the questionnaire. Finally, after each condition, participants were asked to recount *in their own words* what the experience had felt like to them, plus any other remarks they would like to make about the experiment itself. The total experiment took approximately 45 minutes to complete.

3. Results

The rubber hand illusion was measured with an 11 item questionnaire [1] and a drift measure. Furthermore, qualitative data were obtained from the participant's open-ended descriptions. Results from the questionnaire and drift measures will now be presented separately, followed by some illustrative quotes from the participants, recorded after each session.

3.1 Questionnaire

Scores on the questionnaire items for the three experimental conditions are reported in Figure 2. A clear picture emerges of the rubber hand illusion being strongest in the unmediated condition, followed by the VR and lastly the MR one. Similar to the findings by Botvinick and Cohen [1], the first three items showed greatest variance and effects of our manipulations. These were studied more rigorously employing repeated measures analyses of variance (REMANOVA).

The first item ('It seemed as if I were feeling the touch in the location where I saw the rubber hand touched') was analysed in a REMANOVA with Mediation as the independent factor. This factor was significant ($F(2,46) = 10.70$, $p<.001$, partial eta squared $=.32$)[1]. Subsequent contrast analyses revealed a significant difference between the unmediated condition and the two mediated conditions ($p<.001$), but not between the VR and MR condition.

[1] Partial eta squared is an estimate of the degree of association between the dependent and independent variables for the sample and can be interpreted as the proportion of variance in the dependent variable that is attributable to the effect of the independent variable. It is used as an indicator of effect size and its value varies between 0 and 1.

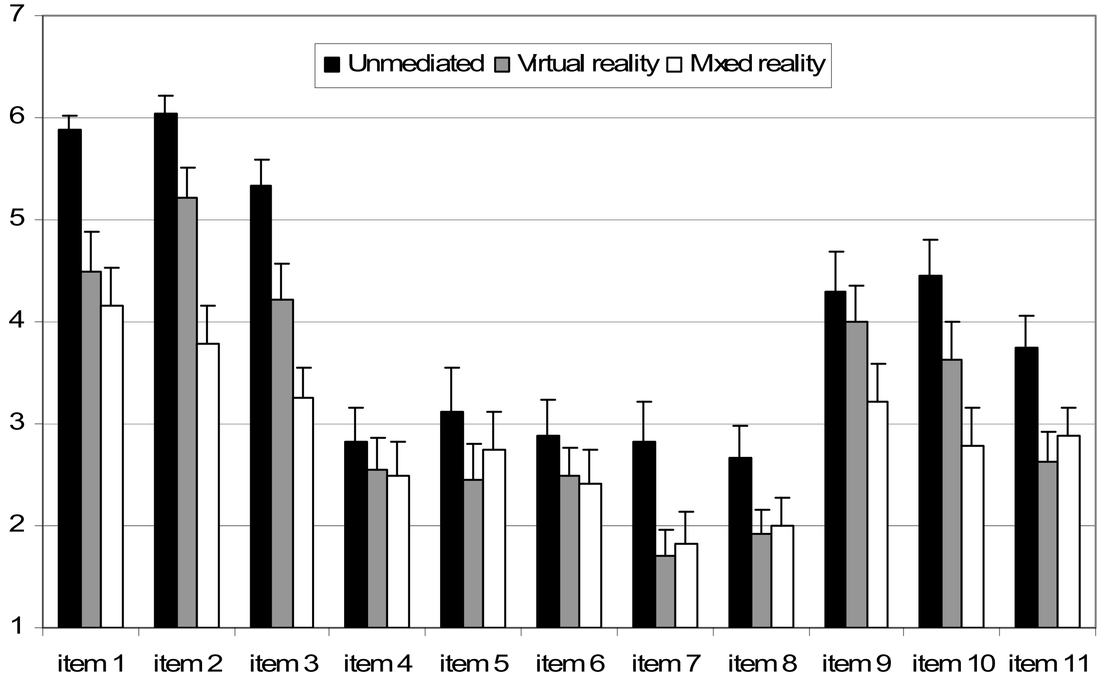

Figure 2: Questionnaire data, presenting means and standard errors of each item for the three experimental conditions. Item 1- It seemed as if I were feeling the touch in the location where I saw the rubber hand touched; item 2 – It seemed as though the touch I felt was caused by the paintbrush touching the rubber hand; item 3 - I felt as if the rubber hand were my hand; item 4 – It felt as if my hand were drifting towards the rubber hand; item 5 – It seemed as if I had more than 1 left hand or arm; item 6 – It seemed as if the touch I was feeling came from somewhere between my own hand and the rubber hand; item 7 – It felt as if my hand was turning rubbery; item 8 – It appeared as if the rubber hand were drifting towards my hand; item 9 – The rubber hand began to resemble my hand in form; item 10 – The rubber hand began to resemble my hand in texture; item 11 – It felt as if my hand was in the rubber hand.

Similar analyses with the second item ('It seemed as though the touch I felt was caused by the paintbrush touching the rubber hand') again revealed a significant effect of Mediation ($F(2,46)=25.87$, $p<.001$, partial eta squared = .53). This time all contrasts were significant ($p<.001$).

Analyses of the third item ('I felt as if the rubber hand were my hand') also resulted in a significant effect of Mediation ($F(2,46)=15.98$, $p<.001$, partial eta squared = .41) and all contrasts significant ($p\leq.01$).

The remaining items showed similar patterns as those described earlier, although in general effects were smaller and not always significant. As a final check, we performed a REMANOVA with the mean score on the 11 items as the dependent variable, and Mediation, Gender, Handedness and Experimental Order as independent variables. Again, Mediation turned out significant ($F(2,9)=8.39$, $p=.009$, partial eta squared = .65), while no remaining significant main or interaction effects emerged.

3.2 Drift

Drift measurements for the three experimental conditions are summarized in Figure 3. Although less clear, the pattern resembles the one found in the questionnaire data: proprioceptive drift of the left hand towards the location of the rubber hand is strongest in the unmediated condition, and weaker in both mediated conditions. A REMANOVA with drift dependent and Mediation independent resulted in marginally significant effects ($F(2,42)=2.64$; $p=.08$). After discarding 1 outlier who had standardised scores over 1.96 in all conditions[2], differences became a bit more pronounced, resulting in a significant effect of Mediation ($F(2,40)=3.71$, $p=0.03$, partial eta squared=.16). Contrast analyses revealed significant differences only between the unmediated condition on the one hand and the mediated conditions on the other.

3.3 Open-ended descriptions

The open-ended description proved to be quite informative. In the unmediated condition there were many cases in which participants were using descriptions that signalled a sense of bodily ownership of the rubber hand. For instance:

"*The feeling seems to build up the first few minutes and then, all of a sudden, the hand feels like my own. And after a while they start to look the same as well!*"

"*Soon you have the feeling the rubber hand is really your hand, you can really feel it being touched.*"

[2] This criterion is known as Grubb's test [11], and is similar to discarding data that differ more than two standard deviations from the mean.

Participants remarked that the illusion was particularly vivid when somewhat more force was applied by the experimenter, and the fingers of the rubber hand moved a little as a result. In both the VR and MR conditions there were several instances where participants also reported a strong sensation of ownership of the mediated rubber hand. In the VR condition, a number of participants also claimed that they felt as if the projection of the rubber hand was a projection of their own hand:

"*I had a feeling I was looking at a projection of my own hand.*"

"*It soon appeared as if the projection was my own hand, and my own hand was being touched.*"

Interestingly, in the MR condition some participants noted that the flat image appeared to obtain volume:

"*It felt as if the projection became three-dimensional, just like my own hand.*"

"*The illusion was not strong, but the image appeared to become 3D as time passed.*"

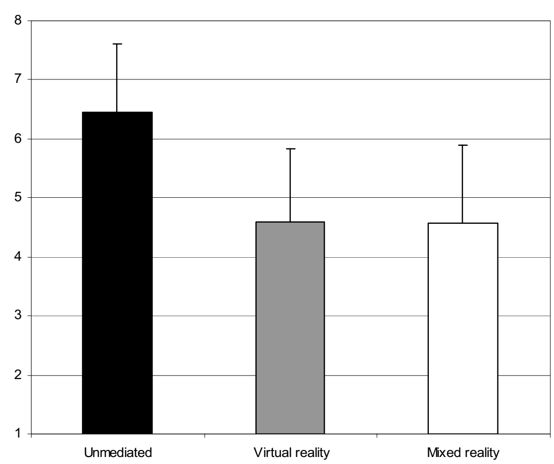

Figure 3: Mean drift in cm for the three experimental conditions. Error bars indicating standard error.

4. Discussion

Our questionnaire results in the unmediated condition clearly replicate the original Botvinick and Cohen results, although with somewhat lower variability in the data. This was to be expected as we selected participants on the basis of a pilot test that showed they were sensitive to the rubber hand illusion to some extent (only 6 of the 30 people tested did not reach this criterion). Nevertheless, the results of the unmediated condition illustrate that the rubber hand illusion can be reliably reproduced when similar procedures are being employed. The existence of the rubber hand illusion demonstrates that intermodal correlations between vision, touch, and proprioception can specify self-attribution of a non-self object [1]. That is, the rubber hand becomes part of the body image, thereby illustrating that the body image is a plastic, temporary construction that can be altered within a relatively short time-span.

The results of the self-report and drift measurements for the mediated conditions indicate that the rubber hand illusion still occurs, albeit to a significantly lesser degree than in the unmediated condition. This result partially contradicts Armel and Ramachandran's [2] claim that the rubber hand illusion is purely the result of Bayesian learning, whereby reliable correlations of visuotactile events are necessary and sufficient by themselves to constitute self-attribution. If this were true, no difference ought to be found between the VR and the unmediated condition, for instance. The fact that we did find a significant difference, however, points to the role of top-down mechanisms that specify requirements for a plausible and congruent (hand-shaped) visual object, if it is to be integrated within the body image. It should be noted however, that Armel and Ramachandran's own results also point to a potential role of top-down mechanisms, as both the subjective ratings and the electrodermal responses were significantly lower in the tabletop condition as compared to the rubber hand condition. Moreover, our results are in agreement with Tsakiris and Haggard [3] who also argue in favour of a combination of bottom-up and top-down processes in explaining the rubber hand illusion. Based on our results, we can argue that the top-down cognitive body representation needs to include a specification of the 3D shape of the hand, as this was the main difference between the projected (VR) and unmediated rubber hand conditions. In the near future, we will employ stereoscopic imaging to further investigate this issue

The VR condition provided participants with a more vivid illusion than the MR condition. This was also in line with our hypothesis, based on the assumption that in the MR condition, like in Armel and Ramachandran's table condition, there was an inconsistency in texture between the felt stimulation on one's skin, and the observed stimulation on the tabletop. This inconsistency was not present in the VR condition. However, after analysing the open-ended descriptions, an alternative explanation for the difference between the VR and MR results also needs to be considered. It appears that a significant number of participants had a quite convincing illusion that the remotely located rubber hand was their own, which was then subsequently being displayed in front of them. None of the participants mentioned this after the MR condition – this would not have made sense as the stimulation was happening on the table in front of them. In the MR condition then, the illusion appeared to suffer somewhat from the conflict between the real brush and the mediated hand. This points to the basic challenge of creating seamless perceptual fusion between the real and the virtual in mixed reality environments. Clearly, in our experiment, this was not yet the case, although for some participants, only in the MR condition, it appeared as though the 2D image became 3D. This illusion could be related to the perceptual system solving the "contradiction" of watching a flat hand being stroked by a 3D brush, and simultaneously feeling one's own unseen hand being stroked.

Overall, our experiment demonstrated that we can produce the rubber hand illusion using media, albeit somewhat less vivid than in the unmediated case. We have shown that form factors play a significant role in the occurrence and vividness of the rubber hand illusion, a fact that contradicts an exclusive adherence to Bayesian principles of statistical correlation. The fact that we can reproduce the rubber hand illusion under mediated conditions is promising for two reasons. First, to obtain a deeper understanding of the form, location, and temporal factors influencing the rubber hand illusion it is necessary to have complete and systematic control over the variables one may want to manipulate. Mediated environments provide such a level of control, combining ecological validity with the ability to systematically tweak relevant variables, and allow for precise replication of conditions [12]. Secondly, the extent to which the mediated rubber hand illusion occurs may in itself provide the research community with an interesting evaluation metric of the quality of the particular media environment under study. The fact that the vividness of the rubber hand illusion varied significantly across conditions in the experiment reported in this paper bodes well for the sensitivity of this measure.

In sum, the same sensorimotor and brain systems responsible for our sense of bodily boundaries are also remarkably adaptable to include non-biological artefacts within the perceptual-motor loop, provided reliable, real-time intersensory correlations can be established, and the artefact can be plausibly mapped onto the body image. When we interact with virtual or remote environments using intuitive interaction devices, isomorphic to our sensorimotor abilities, the real-time, reliable and persistent chain of user action and system feedback will effectively integrate the technology as a phenomenal extension of the self. This fluid integration of technology into the perceptual-motor loop eventually may blur the boundary between our 'unmediated' self and the 'mediating' technology.

Acknowledgements

The authors would like to thank Thomas Beffers, Ferry van Mil, Arjan Tekelenburg, and Bartjan Rijs for carrying out the experiment reported here as part of their Human-Technology Interaction Design Track. We also gratefully acknowledge Martin Boschman for assisting with the laboratory set-up, and Otto Bock Benelux B.V. in Son en Breugel, The Netherlands for donating the prosthetic left hand used in this experiment.

References

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

[2] K.C. Armel, V.S. Ramachandran. Projecting sensations to external objects: Evidence from skin conductance response. *Proceedings of the Royal Society of London B, 270*, 1499-1506, 2003.

[3] M. Tsakiris, P. Haggard. The rubber hand revisited: Visuotactile integration and self-attribution. *Journal of Experimental Psychology: Human Perception and Performance, 31*, 80-91, 2005.

[4] R.B. Welch. The effect of experienced limb identity upon adaptation to simulated displacement of the visual field. *Perception and Psychophysics, 12*, 453-456, 1972.

[5] R.B. Welch, D.H. Warren. Immediate perceptual response to intersensory discrepancy. *Psychological Bulletin, 88*, 638-667, 1980.

[6] R.B. Welch. Adaptation of space perception. In: K.R. Boff, L. Kaufman, and J.P. Thomas (Eds.), *Handbook of Perception and Human Performance*, pp. 24.1-24.44. New York: Wiley, 1986.

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

[8] M. Botvinick. Probing the neural basis of body ownership. *Science, 305*, 782-783, 2004.

[9] W.A. IJsselsteijn. Towards a neuropsychological basis of presence. *Annual Review of CyberTherapy and Telemedicine, 3*, 25-30, 2005.

[10] V.S. Ramachandran, W. Hirstein, D. Rogers-Ramachandran, D. Phantom limbs, body image, and neural plasticity. *Int. Brain. Res. Org. News, 26* (1), 10-11, 1998.

[11] F. Grubbs. Procedures for detecting outlying observations in samples. *Technometrics, 11*, 1–21, 1969.

[12] J. M. Loomis, J.J. Blascovich, A.C. Beall. Immersive virtual environment technology as a basic research tool in psychology. *Behaviour Research Methods, Instruments, & Computers, 31*, 557-564, 1999.