Investigating the Persistence of the Rubber Hand Illusion: The Effects of Shape and Texture

Antal Haans Wijnand IJsselsteijn Yvonne de Kort
Eindhoven University of Technology, The Netherlands
E-mail: A.Haans@tm.tue.nl, W.A.IJsselsteijn@tm.tue.nl, Y.A.W.d.Kort@tm.tue.nl

Abstract

In this paper, we have investigated the extent to which the rubber hand illusion is affected by visual discrepancies between the artificial rubber hand and a real hand. In contrast to earlier studies, we have explored the effects of shape and texture independently, by systematically manipulating these qualities of the artificial object. We found that discrepancies in terms of shape (η_p^2 = 71.7%, $p < .01$) are more detrimental to the illusion than discrepancies in terms of texture (η_p^2 = 16.0%, $p = .05$). Surprisingly, no such differences were found in the extent to which participants misperceived the position of their occluded hand (i.e., proprioceptive drift; $\eta^2 \leq 3.5\%$, $p \geq .38$). Overall, our results corroborate earlier studies that point towards a role of top-down knowledge about the body schema in engendering the illusion.

1. Introduction

"Reality is merely an illusion, albeit a very persistent one."

-Albert Einstein

Historically, the human senses were considered as isolated modalities, each contributing autonomously to perception. Already in the 1830's, several experiments revealed the inadequacy of this view by showing that there is considerable crossmodal transfer between different modalities. In 1838, Johannes Peter Müller [8] realized that the ventriloquist illusion (during which the ventriloquist tricks the audience in believing that the words are coming from the doll's mouth instead of his or her own) was an exception to the autonomous nature of the senses (see e.g., [2]). At about the same time, Brewster [4] showed that people, when looking at a concave object through an optical device that inverted the indentation of the object, would feel the indentation in the same inverted manner when they tactilely explored the object. However, it took more than a century before researchers started to provide theoretical explanations for the crossmodal effects found in these illusions (for an overview see e.g., [2]).

Understanding crossmodal transfer is particularly relevant for the design of multimodal human-system interfaces. Studies on the encoding of peripersonal space and on tool-use (for a review, see [7]), illustrate that we can integrate technological devices as a phenomenal extension of the self, provided that they are matched naturally to our sensorimotor abilities (see [5]). A better understanding of multimodal illusions might enable genuine embodied interaction with technology, perhaps eventually blurring the boundary between our 'unmediated' self and the 'mediating' technology [6].

In the current study, we aim to developed a better understanding of the recently discovered multimodal "*rubber hand illusion*", by investigating the role of top-down knowledge in engendering the illusion.

1.1. The Rubber Hand Illusion

When experiencing the rubber hand illusion, which was first reported by Botvinick and Cohen [3], a person experiences an artificial hand to be his or her own, and it is induced when both the participant's concealed real hand and a fake hand are touched simultaneously and synchronously. When, instead, a small asynchrony is introduced between the stimulation of the participant's occluded hand and that of the fake hand, the illusion does not occur [1,3,12].

An adult human has a long time of experience with his or her body and intuitively the body-image would seem resistant to change. However, the rubber hand illusion illustrates that the right kind of synchronous multimodal stimulation can radically alter our sense of

Proceedings of ENACTIVE05
2nd International Conference on Enactive Interfaces
Genoa, Italy, November 17th-18th, 2005

bodily boundaries, thereby providing evidence for the malleability of the brain in accommodating perceived bodily alterations (e.g., [1,6]).

Next to a sense of ownership, the rubber hand illusion also results in a distortion of position sense [3,6,12]. In other words, after experiencing the illusion, participants incorrectly perceive the location of their occluded arm in the direction of the rubber hand. Moreover, when the artificial hand is "harmed" (e.g., by bending the dummy finger in an anatomically incorrect and thus "painful" manner), participants show clear physiological fear related responses [1].

1.1.1. Underlying perceptual mechanisms. According to Botvinick and Cohen [3], the rubber hand illusion is the result of resolving the ambiguity in the information received from touch and proprioception on the one hand, and from vision on the other. Yet, one might wonder why the experimental situation (i.e., my own hand is touched out of sight in synchrony with a rubber hand being touched) is actually experienced as ambiguous. Another such example is provided by Walton and Spence [13], who state that people find it difficult to selectively attend to a stimulus in one modality and at the same time ignore a synchronous stimulus in another modality. But why exactly is this so difficult?

Perhaps to date, the best answer is provided by the so-called Bayesian multimodal integrations [1]. Bayesian multimodal integration allows the brain to extract correlations between the information received from different modalities upon which it reconstructs a meaningful, and in the above examples erroneous, representation of whatever is out there, including one's body image. In other words, when experiencing the rubber hand illusion, the seen and felt stimulation co-occur with such a high probability, that one's brain cannot do else but decide that the artificial object is part of one's own body. If the brain is not able to make a sufficient correlation between vision and touch, for example when a small asynchrony is introduced between the stimulation of the real hand and the artificial hand, then the rubber hand illusion does not occur (also [12]).

Compared to kinesthesia, vision has a higher spatial acuity (see e.g., [7]), and a second process involved in the rubber hand illusion seems to be visual capture. That is, when the brain is confronted with an incongruence between information received from vision and kinesthesia, it would rely more on vision, recalibrating the position of the occluded arm in the direction of the rubber hand.

1.1.2. Resistance to visual discrepancies. According to Armel and Ramachandran [1], a strong correlation between felt and seen stimulation is sufficient for the rubber hand illusion to occur. They report that people remain susceptible to the illusion even when the rubber hand is placed at a distance of 0.91 meter from the real hand or when the table top is touched instead of an artificial hand (also, [9]). Participants experienced psychological and autonomic arousal (objectively assessed by recording skin conductance response), even when the table-top was "harmed" by pulling a band-aid off the table (note that the experimenters also placed a band-aid on the participant's occluded hand before the start of the experiment). Therefore, the authors conclude that the rubber hand illusion is highly resistant to top-down knowledge about one's body schema; it does not matter how large the visual discrepancies between a real hand and the artificial object, the rubber hand illusion will still occur. However, they do hypothesize that discrepancies in the nature of *expected* felt and seen touch will diminish the illusion. For example, when the table top and the real hand were both touched on the band-aid (i.e., a shared texture) the illusion was more vivid. The authors hypothesize that consistencies between expected and felt touch allow the brain to extract stronger correlations between vision and touch and therefore expect that people will experience a more vivid illusion when a skin-like textured sheet is stimulated, instead of the tabletop. However, one could argue that any consistencies between expected and felt touch depend first of all on visual information about discrepancies in texture.

In contrast, Tsakiris and Haggard [12] and IJsselsteijn, de Kort and Haans [6] argue that Bayesian multimodal integration of vision and touch by means of simultaneous and synchronous stimulation are, although necessary, not sufficient for the rubber hand illusion to occur. Tsakiris and Haggard showed that when the rubber hand was placed in an orthogonal position relative to the participant's occluded hand, or when a wooden stick was used as the artificial object, participants did not show a significant re-calibration of position sense. Similarly, IJsselsteijn et al. showed that participants were significantly less susceptible to the illusion, when they looked at a two-dimensional projection of the rubber hand, compared to situations where the rubber hand was physically present on the table. The results of these latter two studies seem to indicate that top-down knowledge about the body schema has a considerable effect on the rubber hand illusion, and that at least some correspondence between the artificial object and the human body is required.

Unfortunately, the studies of Armel and Ramachandran [1] and Tsakiris and Haggard [12] are difficult to compare as they use different measures to assess the experienced strength of the illusion. Secondly, published studies do not allow to make independent comparisons between the effect of discrepancies in terms of shape and texture, as these qualities have not been systematically manipulated. For example, Tsakiris and Haggard [12] used a wooden stick as the artificial object, which differs in both shape and texture from the participants' own hand.

1.1.3. Measuring the strength of the illusion. Most of the published studies on the rubber hand illusion rely on self-report and proprioceptive drift measures to assess the strength in which people experienced the illusion.

IJsselsteijn et al. [6] adapted and extended the self-report measure of Botvinick and Cohen [3]. The original version of the measure contained nine 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". IJsselsteijn et al. divided 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" into two separate items, one for resemblance in terms of shape, and the other for the resemblance in terms of texture. Secondly, they added one item which described a sensation that a number of people reported during the pilot phase of their study: "It felt as if my hand was inside the rubber hand". Participants used a seven point response format, running from "not at all" to "completely", to indicate the extent to which each of the 11 statements matched their own experiences. By contrast, Armel and Ramachandran [1] used a single item measure, asking people to rate the extent to which the rubber hand felt as their own. Participants could respond by means of a 10-point response format.

Proprioceptive drift represents the extent to which people misperceive the position of their occluded hand and is measured either by having them point the felt location of the occluded hand with their other hand (e.g., [3,6]), or by having them verbally report on the felt position using a ruler [12]. Since people differ in the initial felt position of their occluded hand (i.e., before experiencing the illusion), most studies (e.g., [6,12]) use change or difference scores (i.e., post-exposure drift minus pre-exposure drift) to assess the effect of the rubber hand illusion on position sense. Unfortunately, the authors of these studies report on these difference scores only, leaving the reader to speculate about the observed size of pre-exposure drift and, thus, about the exact size of the effect of the illusion on position sense.

1.2. Aim of the present study

In the present study, we will investigate the extent to which the rubber hand illusion is affected by visual discrepancies between the artificial object and a real hand. In contrast to earlier studies (e.g., [12]), we explore the effects of shape and texture independently, by systematically manipulating these qualities of the artificial object. We used a self-report measure as well as a proprioceptive drift measure to asses the experienced strength of the illusion, and to allow for a more direct comparison with the results reported by Armel and Ramachandran [3] and Tsakiris and Haggard [12].

We expect to corroborate the findings of Tsakiris and Haggard [12] and IJsselsteijn et al. [6], whose results indicate that top-down knowledge about the body schema has a considerable effect on the rubber hand illusion, and that at least some correspondence between the artificial object and the human body is required. Secondly, we want to test whether similarities in terms of shape or similarities in terms of texture are the most important in engendering the illusion.

2. Methods

2.1. Design

A two (shape and no shape) by two (natural and non-natural skin texture) repeated measures experiment was conducted. The experiment consisted of four sessions. In each of these sessions a different artificial object was used. For the *hand Shape with natural Texture condition* (abbreviated as ST) a cosmetic prosthesis of a man's left hand was used, which was highly realistic in terms of skin texture, color and shape. For the *hand Shape with no natural Texture condition* (S¬T) a white latex glove was fitted over the cosmetic prosthesis to modify texture and color, but not shape. For the *no hand Shape with natural Texture condition* (¬ST) a flat sheet (size 24 X 13 cm) of the same material as the cosmetic prosthesis was placed in front of the participant (as suggested by [1]). Finally, for the *no hand Shape with no natural Texture condition* (¬S¬T) no object was used, leaving only the white colored table top, replicating one of Armel and Ramachandran's [1] conditions. The ST condition was always in the first session as this condition was used to select only those participant who were susceptible to the original version of the illusion. The order of the remaining three conditions was balanced across participants over the remaining three sessions.

2.2. Participants

The present sample was drawn from employees and students of the Human-Technology Interaction department at Eindhoven University of Technology. Twenty-six persons participated in the experiment. All were tested on their susceptibility to experience the rubber hand illusion at the beginning of the experiment. Three out of 26 (i.e., 11.5%) did not experience the original version of the illusion (i.e., ST condition) and were excluded from participation in the experiment. Of the remaining 23 participants, the mean age was 22.3 (*SD* = 2.2; range 18 to 27 years); 14 were male; 18 were right handed. All students received a standard compensation of 7 € for their participation.

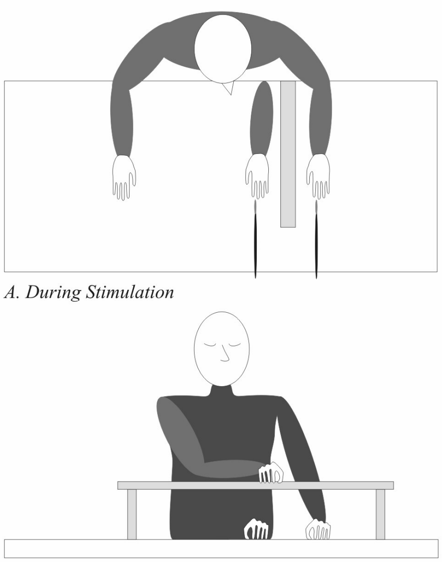

A. During Stimulation

B. Measuring Proprioceptive Drift

Figure 1. The Experimental Setup

Note, that Figure A shows the setup during the stimulation of the participant's hand and the artificial object (top view) and that Figure B shows the setup during the measurement of proprioceptive drift (front view).

2.3. Experimental Setup

The participants sat upright at a table opposite to the experimenter (see Fig. 1A). The experimenter ensured that the participant was in a comfortable position with both arms resting comfortably on the table with the palms of the hands down. Before the beginning of the experiment, the participants were asked to put on a jacket. The left arm sleeve had been removed and was placed on the table in front of the participant, in a way that suggested that the artificial object belonged to the participant's body. The distance between the participant's left hand and the artificial object was 30 cm and the same distance was used in each session.

When the participants were ready to begin, the experimenter asked them to put their left hand on a cross on the table (i.e., 30 cm from the artificial object). The participants were instructed to keep their left hand motionless during the whole session. Subsequently, the experimenter placed a wooden screen between the participant's left hand and the artificial object to occlude the hand from view. Next, the experimenter used two small brushes to synchronously stroke congruent positions on both the artificial object and the participant's occluded hand for five minutes. This procedure was repeated for each of the four sessions. After each session, several measures were used to assess the strength in which the participants experienced the illusion.

2.4. Measures

We used two different measures to assess the strength in which the participants experienced the illusion, namely a self-report and a proprioceptive drift measure. Participants were also encouraged to write down their experiences after each session in their own words.

2.4.1. Self-report measure. The questionnaire was adopted from IJsselsteijn, et al. [6] (as described in Section 1.1.3). The measure consisted of 11 statements describing specific perceptual effects associated with the rubber hand illusion. The term "rubber hand" in the items was replaced by the appropriate term for each condition (e.g., "rubber sheet"). Participants were asked to indicate the extent to which each statement matched their own experiences. A seven point response scale was used for each item, ranging from 0 (not at all) to 6 (completely). Participants were asked to complete the questionnaire after each session. Across all sessions there were no missing responses. Scores were calculated by taking the mean of each person's responses to the 11 items. The internal consistency of each of the four scales, was $\alpha = .62$ for the rubber hand (ST) condition, $\alpha = .81$ for the latex hand (S¬T) condition, $\alpha = .85$ for the rubber sheet (¬ST) condition, and $\alpha = .90$ for the table top (¬S¬T) condition.

2.4.2. Proprioceptive drift. Proprioceptive drift was measured at the end of each 5 minute session. At that time, the experimenter asked the participants to close their eyes and reminded them to keep their left hand in place on the table. Next, the experimenter removed the hardboard screen and put a small table (30 by 80 cm with a height of 24 cm) over the participant's left hand and the artificial object (see Figure 1B). After the table was in place, the experimenter took the participant's right hand and placed it at the beginning of the table. Next, the participants were instructed to indicate the position of their left hand by moving their right hand over the table until they felt that both hands were aligned. Drift toward the participant's right hand (i.e., toward the rubber hand) was coded as positive. After the experimenter had marked the indicated position and had removed the table, the participants were allowed to open their eyes and move their left hand. As a baseline measure for comparison, we also measured a person's pre-exposure proprioceptive drift. This pre-exposure drift was measured at the beginning of the experiment by taking the mean of three successive trials (internal consistency $\alpha = .95$).

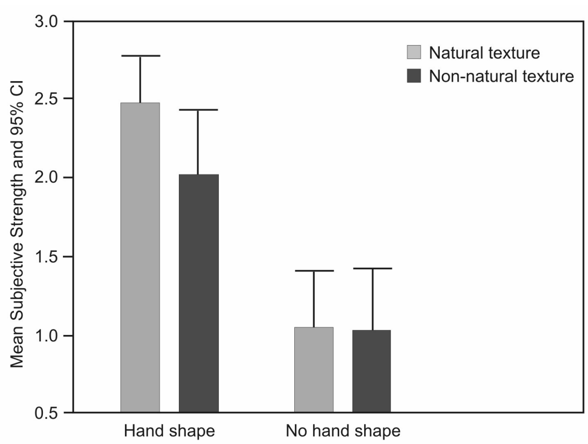

Figure 2. Self-Report Measure

3. Results

Our findings are reported in three sections. First, we present the results of the analyses of the self-report measure. Second, we present the analyses of the proprioceptive drift measure. Third, we will present some of the free format remarks that were made by the participants after each session.

3.1. Subjective measure

The mean experienced strength of the illusion and their corresponding 95% confidence intervals are depicted in Figure 2 for each experimental condition (see also Table 1). Results indicate a considerable main effect of Shape (see, row effect in Table 1A), and a small main effect of Texture (see, column effect in Table 1A). By examining the residuals after the main effects have been removed (see e.g., [10,11]), we find that there is also a non-zero, yet small, interaction effect between Shape and Texture (see, the Residuals in Table 1C).

Table 1. Means, Standard Error and Contrasts

A. Means	Texture			
Shape	Natural	Non-natural	Row mean	Row effect
Hand	2.47 (.15)	2.02 (.20)	2.25	+0.60
No hand	1.05 (.17)	1.04 (.19)	1.05	-0.60
Column mean	1.76	1.53	**1.65**	
Column effect	+0.12	-0.12		
B. Estimates				
Hand	2.36	2.13		
No hand	1.16	0.93		
C. Residuals				
Hand	+0.11	-0.11		
No hand	-0.11	+0.11		

Note, that figures in parenthesis represent standard errors; the figure in **bold** represents the overall or grand mean. Row and column effects (i.e. main effects of shape and texture) are calculated by subtracting the grand mean from the row mean and column mean respectively. The estimates are calculated by summing the grand mean and column and row effects for each cell. Residuals (i.e., the interaction effect) are calculated by subtracting the cell estimates from the original cell means.

To calculate the significance of each of the three effects, we performed a two by two repeated measures ANOVA with the experienced strength of the illusion as the dependent variable. We found a significant main effect for Shape [$F(1,22) = 55.6$, $p < .01$, and Partial Eta Squared $\eta_p^2 = 71.7\%$] and Texture [$F(1,22) = 4.2$, $p = .05$, $\eta_p^2 = 16.0\%$] on the experienced strength of the illusion. In addition, the interaction effect between Shape and Texture was also found to be significant [$F(1,22) = 4.4$, $p = .05$, $\eta_p^2 = 16.6\%$].

3.2. Proprioceptive drift

Participants showed, on average, a pre-exposure drift of $M = 3.5$ cm with $SE = 0.9$ (i.e., towards the rubber hand). The average proprioceptive drift after the ST ($M = 7.1$; $SE = 1.0$), S¬T ($M = 7.2$; $SE = 1.3$), ¬ST ($M = 6.5$; $SE = 0.89$) and ¬S¬T ($M = 7.4$; $SE = 1.4$) conditions and their corresponding 95% confidence intervals are depicted in Figure 3. We conducted a series of paired sample t-tests to test for differences between pre-exposure (i.e., baseline) drift and the drift after each experimental condition. The proprioceptive drift after the illusion was significantly different from pre-exposure drift for all experimental conditions [$t(22) > 2.7$, $p \leq .02$].

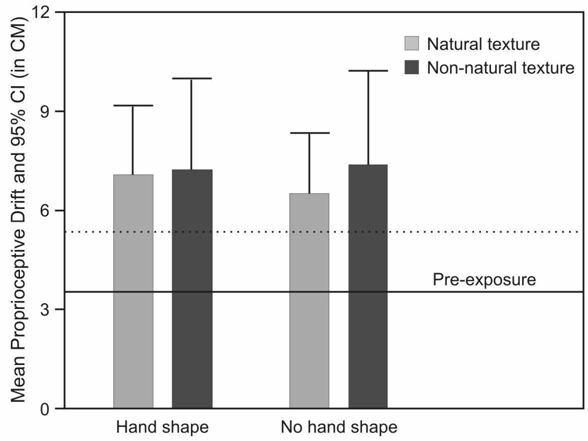

Figure 3. Proprioceptive Drift

Note, that the solid horizontal line represents mean pre-exposure (i.e., baseline) drift and the dotted line the corresponding upper 95% confidence interval.

Secondly, we performed a two (hand Shape, no hand Shape) by two (natural skin Texture, non-natural skin Texture) repeated measures ANOVA with post-exposure Proprioceptive Drift as the dependent variable. We found no significant main effects for Shape [F(1,22) = .1, p = .73, η_p^2 = 0.6%] and Texture [F(1,22) = .8, p = .38, η_p^2 = 3.5%], nor a significant interaction effect [F(1,22) = .4, p = .55, η_p^2 = 1.7%]. Note, that using difference scores as the independent variable (i.e., pre-exposure drift minus post-exposure drift; see Section 1.1.3) has no practical use in a complete within subject design.

3.3. Free Responses

Our participants' descriptions reflected how people commonly experience the illusion. For example, most of our participants mentioned experiencing something strange or mentioned a tingling sensation in their left hand, both of which are commonly encountered in the literature on the rubber hand illusion.

During the S¬T condition, in which a white latex glove was fitted over the cosmetic prosthesis, some participants reported feeling as if their hand was inside the glove, others however reported that they were distracted by the wrinkles in the glove.

More importantly, approximately half of the participants indicated that they did not experience the rubber hand illusion in one or both of the ¬ST (i.e., skin-like sheet) and the ¬S¬T (i.e., tabletop) conditions. Some participants, indicated that they were distracted by the rectangle shape of the skin-like sheet and that they experienced a more vivid illusion in the ¬S¬T condition as opposed to the ¬ST condition.

4. Discussion

In the present study, we have explored the effects of visual discrepancies between the artificial object a and real hand on the strength in which people experience the rubber hand illusion. In contrast to earlier studies (e.g., 12), we have explored the effects of shape and texture independently, by systematically manipulating these qualities of the artificial object. The experienced strength of the illusion was assessed by two measures, namely a self-report measure and a proprioceptive drift measure.

On the 11-item self-report measure, we found a considerable effect of shape (η_p^2 = 71.7%, p < .01), indicating that the participants experienced a stronger illusion when the artificial object resembled their own hand in terms of shape. By contrast, the effect of texture was relatively small and in favor of the texture that resembled the human skin (η_p^2 = 16.0%, p = .05), providing corroborating evidence for Armel and Ramachandran's [1] hypothesis that a skin-like sheet on the tabletop has a positive effect on the rubber hand illusion. Yet, *in contrast* to what we would expect from Armel and Ramachandran's anecdotal observations, we found that visual discrepancies in terms of shape are more detrimental to the illusion than discrepancies in terms of texture, at least for the experimental manipulations used in the present study. Had we manipulated shape and texture differently, we might have found different results. Nevertheless, the manipulations used in the present study were similar to those used or proposed by Armel and Ramachandran.

We also found a small interaction effect between shape and texture (η_p^2 = 16.6%, p = .05) on the self-report measure, indicating that opposite to a non-natural texture, and in addition to the main effect of texture, a natural texture applied to a non hand-shaped object actually had a small detrimental effect on the experienced strength of the illusion (see Table 1C). A preliminary explanation for this interaction effect could be that some participants were distracted by the rectangle shape of the skin-like sheet and therefore experienced a stronger illusion in the tabletop condition, despite its non-natural texture (see Section 3.3).

In contrast to our findings on the self-report measure, we found no significant effects of shape (η_p^2 = 0.6%, p = .73) and texture (η_p^2 = 3.5%, p = .38) on the proprioceptive drift measure. Unexpectedly, participants showed the same amount of proprioceptive drift, irrespective of whether the artificial object resembled a real hand or not. In contrast to our findings, Tsakiris and Haggard [12] report a significant decrease in drift when a wooden stick was used as the artificial object instead of a rubber hand. Our method of measuring proprioceptive drift differed considerably from that used by Tsakiris and Haggard, and perhaps their method was less prone to measurement error and thus more sensitive to differences in drift.

However, we did find significant differences between pre-exposure (i.e., baseline drift) and post-exposure drift for all combinations of shape and texture [$t(22) > 2.7$, $p \leq .02$], indicating that people experience the illusion even when the artificial object does not resemble the human hand. By contrast, Tsakiris and Haggard [12] found that proprioceptive drift did not exceed its baseline value when a wooden stick was used. Although our proprioceptive drift measure indicates that a resemblance between the artificial object and a real hand is not necessary for the illusion to occur, half of our participants reported not to experience the illusion in the ¬ST (i.e., skin-like sheet) or ¬S¬T (i.e., tabletop) condition (see Section 3.3).

On the one hand, our self-report measure and the free responses show that visual discrepancies in shape and texture are detrimental to the rubber hand illusion. On the other hand, our proprioceptive drift measure shows that this is not necessarily the case. For now, we can only speculate about the exact origins of this unexpected finding. Published research (e.g., [14]) on proprioceptive drift shows that a five minute period without visual feedback about the position of the stationary arm, in itself, is sufficient for people to show comparable distortions of position sense. Perhaps, proprioceptive drift is not a valid measure to assess the experienced strength of the rubber hand illusion. Regrettably, the study of Tsakiris and Haggard [12] relies exclusively on drift and no free response or any other type of subjective reports are used. More research is necessary to conclusively determine the effect of the rubber hand illusion on proprioceptive drift.

Despite the abovementioned limitations, we believe that we have corroborated the findings of Tsakiris and Haggard [12] and IJsselsteijn et al. [6] by showing that the rubber hand illusion is not resistant to top-down knowledge about the body-schema; the larger the resemblance between the artificial object and the body schema (especially in terms of shape), the stronger the experienced rubber hand illusion. Yet, interpersonal differences seem to play a mayor role as well. People differ considerably and consistently in their ability to experience the illusion; some people experience the illusion even during the tabletop condition, while others are resistant to the illusion, even when a highly realistic cosmetic hand prosthesis is used (in this study three out of 26 participants). Possibly, the experienced strength of the illusion depends on trait-like interpersonal differences in susceptibility as well as on the degree in which the artificial object resembles the body schema. Perhaps, Armel and Ramachandran's [1] conclusion that the rubber hand illusion occurs irrespective of the visual qualities of the artificial object, is based on a sample of highly susceptible people. Therefore, future studies should definitely take these interpersonal differences into account.

Allowing the brain to make correlations between vision and touch by means of simultaneous and synchronous stimulation is on itself not sufficient to engender the rubber hand illusion, as top-down knowledge of the body schema prevents most people to surpass the boundaries between their bodily self and the inanimate material world.

Acknowledgement

This research was supported by the JF Schouten School for User-System Interaction Research at Eindhoven University of Technology, Eindhoven, The Netherlands. We gratefully acknowledge Otto Bock Benelux B.V. and Jongenengel Orthopedisch Centrum for kindly donating and manufacturing the materials used in the present study. We also thank Roderick Hoek, Eline Katz, Jerry Maasen, and Ralf van der Ven for their assistance in conducting the experiment, and Florian Kaiser for his comments on earlier drafts of this paper.

References

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

[2] Bertelson, P., & de Gelder, B. (2004). The psychology of multimodal perception. In C. Spence and J. Driver (Eds.), *Crossmodal space and crossmodal attention* (pp. 151-177). Oxford: Oxford University Press.

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

[4] Brewster, D. (1832). *Letters on natural magic.* London: John Murray.

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

[6] IJsselsteijn, W.A., de Kort, Y.A.W., & Haans, A. (2005). Is this my hand I see before me? The rubber hand illusion in reality, virtual reality, and mixed reality. In M. Slater (Ed.), *Proceedings of the 8th International Workshop on Presence* (pp. 41-47). London: University College London.

[7] Maravita, A., Spence, C., & Driver, J. (2003). Multisensory integration and the body schema: close to hand and within reach. *Current Biology, 13,* 531-539.

[8] Müller, J.P. (1838). *Handbuch der physiologie des menschen* (Vol. 2) [Handbook of Human Physiology]. Coblenz, Germany: Hölscher.

[9] Ramachandran, V.S., Hirstein, W., & Rogers-Ramachandran, D. (1998). Phantom limbs, body image, and neural plasticity. *International Brain Research Organization News, 26(1),* 10-11.

[10] Rosnow, R.L., & Rosenthal, R. (1989). Definition and interpretation of interaction effects. *Psychological Bulletin, 105,* 143-146.

[11] Rosnow, R.L., & Rosenthal, R. (1995). "SOME THINGS YOU LEARN AREN'T SO": Cohen's paradox, Asch's paradigm, and the interpretation of interaction. *Psychological Science 6,* 3-9.

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

[13] Walton, M., & Spence, C. (2004). Cross-modal congruency and visual capture in a visual elevation-discrimination task. *Experimental Brain Research, 154,* 113-120.

[14] Wann, J.P., & Ibrahim, S.F. (1992). Does limb proprioception drift? *Experimental Brain Research, 91,* 162-166.