Journal of Experimental Psychology:
Human Perception and Performance
2005, Vol. 31, No. 1, 80–91

0096-1523/05/$12.00 DOI: 10.1037/0096-1523.31.1.80

The Rubber Hand Illusion Revisited: Visuotactile Integration and Self-Attribution

Manos Tsakiris and Patrick Haggard
University College London

Watching a rubber hand being stroked, while one's own unseen hand is synchronously stroked, may cause the rubber hand to be attributed to one's own body, to "feel like it's my hand." A behavioral measure of the rubber hand illusion (RHI) is a drift of the perceived position of one's own hand toward the rubber hand. The authors investigated (a) the influence of general body scheme representations on the RHI in Experiments 1 and 2 and (b) the necessary conditions of visuotactile stimulation underlying the RHI in Experiments 3 and 4. Overall, the results suggest that at the level of the process underlying the build up of the RHI, bottom-up processes of visuotactile correlation drive the illusion as a necessary, but not sufficient, condition. Conversely, at the level of the phenomenological content, the illusion is modulated by top-down influences originating from the representation of one's own body.

Watching a rubber hand being stroked synchronously with one's own unseen hand causes the rubber hand to be attributed to one's own body, to "feel like it's my hand." Attribution can be measured quantitatively as a drift of the perceived position of one's own hand toward the rubber hand. In one study, after 30 min of synchronous stimulation on the rubber hand and the participant's hand, participants mislocated the perceived position of their own hand (Botvinick & Cohen, 1998). Participants judged the position of their hand to be closer to the rubber hand, as if their hand had drifted toward the fake hand. Botvinick and Cohen suggested that the rubber hand illusion (RHI) reflected a three-way interaction between vision, touch, and proprioception: Vision captured touch, resulting in a mislocalization of the tactile percept toward the spatial location of the visual percept. This visual–tactile correlation influenced the felt position of one's own hand.

The end result of the RHI, namely, a visual adaptation of proprioceptive position, is somewhat similar to prism adaptation (Welch, Widawski, Harrington, & Warren, 1979). However, the key feature of the RHI is that it results from correlated visual and tactile inputs. This interplay between vision and touch has been at the center of research on multisensory integration (for a review, see Maravita, Spence, & Driver, 2003). It has been consistently shown that vision usually plays a dominant role over touch and proprioception (Ernst & Banks, 2002). For example, even noninformative vision of body parts exerts a top-down influence on tactile percepts by improving the spatial resolution of touch (Haggard, Taylor-Clarke, & Kennett, 2003; Kennett, Taylor-Clarke, & Haggard, 2001).

The RHI can be used as a method of investigating not only multisensory integration and the interplay between vision, touch, and proprioception, but also the way we perceive our bodily self. The necessary condition for the inducement of the illusion is the presence of synchronized and spatially congruent visual and tactile stimulation. At the phenomenological level, the illusory experience seems like a form of incorporation of a foreign object, as if the rubber hand becomes a *bodily auxiliary* (Merleau-Ponty, 1945/1962), possibly in similar ways to the incorporation of tools (Iriki, Tanaka, & Iwamura, 1996).

Introspective evidence from the original experiment (Botvinick & Cohen, 1998) suggested that participants felt not only as if they were feeling the touch at the location where the rubber hand was seen to be touched but also as if the rubber hand was their own hand. In a sense, their tactile sensations were projected onto the rubber hand, which eventually felt like part of their own body. These observations might reflect the involvement of two separate components. First, there is a bottom-up process of integrating synchronized visual and tactile percepts, which is a necessary condition for producing the RHI. Second, this process produces persistent, vivid phenomenological changes in body representation, namely, the experience that the rubber hand is part of one's own body. Moreover, the content of the changed body representation might be quite different from, and goes beyond, the perception of correlated visual and tactile stimulation.

These observations suggest that the RHI could involve an interaction between general body scheme representations and localized visuotactile integration. Two issues are in need of clarification. First, how do correlated visual and tactile stimuli come to influence the perceived position of the hand? Second, why do correlated visual and tactile stimuli produce this strange phenomenal experience of ownership, that is, do normal participants who know that this is a fake hand feel as if the rubber hand belongs to their body?

We report four experiments. In Experiments 1 and 2, we were principally interested in investigating the modulation exerted by general body scheme representations on the inducement of the illusion. To that extent, in Experiment 1 we manipulated (a) the posture of the viewed rubber hand and (b) the identity of the viewed object, and in Experiment 2 we manipulated the *felt* and

---

Manos Tsakiris and Patrick Haggard, Institute of Cognitive Neuroscience and Department of Psychology, University College London, London, United Kingdom.

Correspondence concerning this article should be addressed to Manos Tsakiris, Institute of Cognitive Neuroscience, Alexandra House, 17 Queen Square, London WC1N 3AR, United Kingdom. E-mail: e.tsakiris@ucl.ac.uk

seen identities of the hands. In Experiments 3 and 4, we focused on the perceptual process of multisensory integration that builds up to the inducement of the illusion. By stimulating one or two fingers, we attempted to see whether the mislocalization of the position of the hand is effected in a global or a local way.

Experiment 1

Experimental Design

The aim of Experiment 1 was to investigate the influence of the viewed stimulus content on the RHI by manipulating first the body configuration, and second, the identity of the viewed object. This experiment was informed by previous suggestions in the literature. For example, Pavani, Spence, and Driver (2000) showed that interference effects of viewing rubber hands, while participants were tapped on their hands, were operative only when the posture of the rubber hand was congruent with a participant's own posture. In that study, participants had to discriminate vibrotactile stimuli delivered to either the index finger or the thumb of their own unseen hand, while they were looking at a rubber hand "holding" distractor lights. The effect of the distractor lights was increased only when the rubber hand was in a congruent position with respect to the participant's hand and was decreased when the rubber hand was in an incongruent position.

Therefore, we first manipulated the posture of the rubber hand. Participants watched a rubber hand in a congruent position with their own hand (0°), or a rubber hand in an incongruent position, rotated by −90° with respect to their own hand.

Our second manipulation of viewing a neutral object represents an attempt to clarify whether the RHI is caused by a purely bottom-up association of visual and tactile events (see Armel & Ramachandran, 2003; Ramachandran & Hirtsein, 1998). This hypothesis suggests that any visual stimulus will induce a similar illusion, provided it is appropriately correlated with tactile stimulation. Therefore, in this condition, the rubber hand was replaced by a wooden stick.

For the RHI to occur, visual and tactile percepts should be temporally synchronized (Botvinick & Cohen, 1998). We therefore used a repeated measures design, in which we compared an experimental condition of synchronous stimulation with a control condition of asynchronous stimulation. In previous studies (e.g., Botvinick & Cohen, 1998), the mode of stroking was a between-subjects factor (see General Discussion).

To recap, the experimental factors were: (a) the viewed object (rubber hand at 0°, rubber hand at −90°, wooden stick) and (b) the mode of stroking (synchronous vs. asynchronous). The conditions were blocked, and each participant performed the conditions in a different random order.

Method and Participants

Figure 1 shows the experimental setup. Participants sat in front of a table. At the beginning of each block, the experimenter placed the participant's left hand at a fixed point inside a frame, the top side of which was covered by one-way and two-way mirrors. The two-way mirror was used to make the rubber hand appear (during stimulation) and disappear (during judgment). At the beginning of each block, both the participant's left hand and the rubber hand were out of sight. A pretest baseline estimate of finger position was obtained prior to stimulation. Participants saw a ruler reflected on the mirror. The ruler was positioned 18 cm above the mirror, to appear at the same gaze depth as the rubber hand. Participants were asked, "Where is your index finger?" and in response, they verbally reported a number on the ruler. They were instructed to judge the position of their finger by projecting a parasagittal line from the center of their fingertip to the ruler. During the judgments, there was no tactile stimulation, and the lights under the two-way mirror were switched off to make the rubber hand and stick invisible, leaving only the ruler visible.

After the judgment, the ruler was removed, and the lights under the two-way mirror were turned on to make the rubber hand (or stick) appear, aligned with the participant's midline. The participants were viewing the rubber hand (or stick) in the same depth plane as their own hand. The distance between the real hand and the viewed object was 17.5 cm.

Stimulation was delivered manually over 4 min by the experimenter with the use of two identical paintbrushes. Participants were always stimulated horizontally on the index finger of their left hand, and the rubber hand was stimulated on the index finger in the same way. Both participants and rubber hands wore identical elastic gloves to eliminate visual, tactile, and auditory differences. The tip of the wooden stick was also covered with the same elastic material for the same reasons. In the synchronous conditions, both the participant's stimulated hand and the rubber hand were stroked simultaneously and at the same location. Stimulation was delivered manually along the index finger from the knuckle to the fingertip. Each stroke lasted approximately 500–1,000 ms. The experimenter immediately repositioned the stick at the knuckle and began the next stroke some 500–1,000 ms after the end of the previous stroke. In the asynchronous conditions,

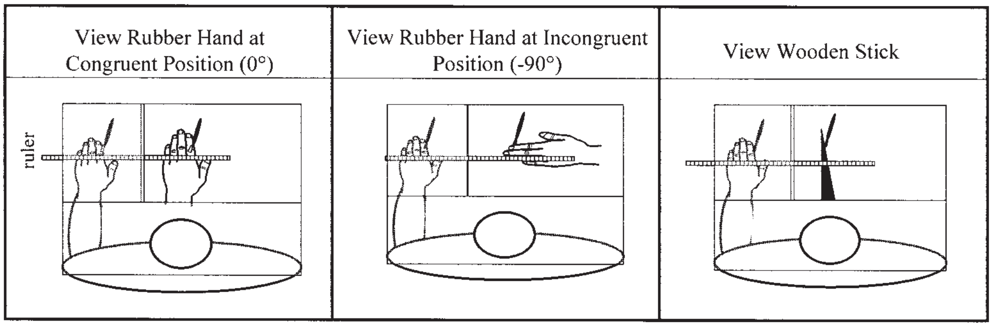

Figure 1. Experimental setup in Experiment 1. Participants saw in different conditions (a) a rubber hand in a congruent posture, (b) a rubber hand in an incongruent posture, or (c) a wooden stick. The participant's left hand was out of sight for the whole duration of the experiment. The rubber hand or the wooden stick appeared, aligned with the participant's midline, only during the stimulation and disappeared during the judgment period.

visual stimulation preceded tactile stimulation, and the asynchrony was randomly varied between 500–1,000 ms.

After the stimulation period, the lights were automatically turned off. The ruler was always presented with a random offset to ensure that participants judged finger position anew on each trial and that they could not simply repeat previous responses. Participants were asked, "Where is your index finger?" After their answer, the ruler was removed, and they were asked to move their left hand and have a rest for a few moments. Following the rest period, their left hand was again passively placed at a predetermined point, under the frame and out of sight. The same process was followed for each condition.

Eight volunteers (mean age = 26.7 years, 4 female and 4 male), all right-handed, with normal or corrected-to-normal vision, participated in this study after giving their informed consent.

Results

We obtained a baseline pretest judgment prior to stimulation and a posttest judgment after stimulation. Judgment errors were calculated as the difference between the perceived and the real position of the participants' fingers. A positive error represents a mislocalization toward the rubber hand. The pretest judgment errors were subtracted from the posttest judgment errors prior to analysis. Figure 2, therefore, shows the change in the perceived position of the hand between the start and end of the stimulation period, across conditions. We use the term *proprioceptive drift* to refer to this quantity.

We first performed a 2 × 3 analysis of variance (ANOVA) on the proprioceptive drifts with two within-subjects factors. The factors were (a) the mode of stroking (synchronous vs. asynchronous) and (b) the viewed object (rubber at 0°, rubber at −90°, wooden stick). No main effects were significant, $F(1, 7) = 0.76$, $p > .05$, for the mode of stroking, and $F(2, 14) = 0.59$, $p > .05$, for the viewed object. The interaction of the two factors was significant, $F(2, 14) = 7.96$, $p < .05$. To investigate this interaction further, we then used simple effects analysis (Howell, 1997) to compare the proprioceptive drift between synchronous and asynchronous conditions for each visual stimulus condition.

Differences between synchronous and asynchronous conditions were significant only when participants saw a rubber hand at a congruent posture with their own hand, $t(7) = 4.25$, $p < .006$, two-tailed. Differences between synchronous and asynchronous conditions when a participant was looking at a rubber hand in an incongruent posture, $t(7) = -1.1$, $p > .05$, two-tailed, or when looking at a wooden stick, $t(7) = -2.04$, $p > .05$, two-tailed, were not significant. Indeed, in these last two cases the drifts for the asynchronous condition were larger than the drifts for the synchronous condition, contrary to prediction.

Inspection of Figure 2 shows a particularly strong drift of the felt position of the hand toward the rubber hand following synchronous stimulation, especially when participants viewed a rubber hand at a congruent posture. However, some drift is seen even after asynchronous stroking for the incongruent posture and stick conditions. This suggests that drift toward the viewed object may reflect a nonspecific effect of visual or tactile stimulation on proprioceptive representation, but may occur due to other factors also. Even during the pretest judgments, participants perceived their finger to be closer to the midline than it really was. This effect might reflect the reported bias in perceiving the position of the hand as closer to the midline than it truly is (Ghilardi, Gordon, & Ghez, 1995). Our interest here is on that part of the positional drift due to visual–tactile integration. This integration component can be defined as the increase in positional drift when visual and tactile stimulation are correlated (i.e., synchronous conditions), over and

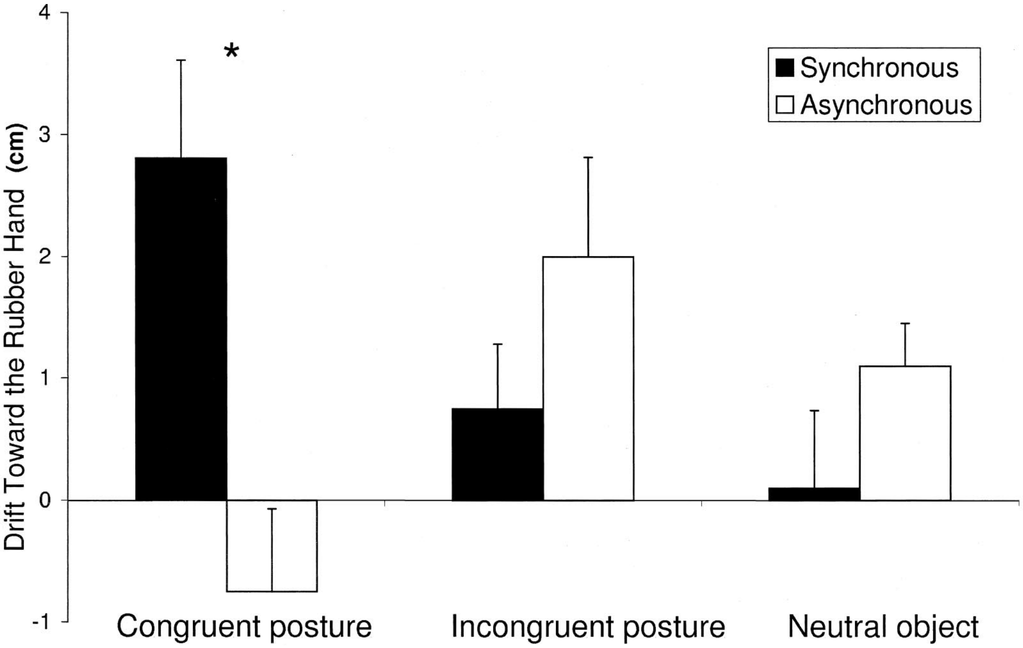

Figure 2. Mean proprioceptive drifts toward the rubber hand in Experiment 1. Error bars indicate standard errors. Asterisk indicates significant differences between synchronous and asynchronous stimulation. Zero represents the felt position of the participant's hand prior to stimulation.

above the drift caused by the same stimuli when they are not correlated (i.e., asynchronous conditions). Therefore, to obtain a more specific measure of the RHI, we subtracted the proprioceptive drifts obtained in the asynchronous conditions from the proprioceptive drifts obtained in the synchronous conditions. We use the term perceptual shifts to refer to this quantity. The perceptual shift measure was positive when participants viewed a congruent hand posture, but it was negative when participants viewed an incongruent posture or a neutral object.

We performed two planned comparisons on the perceptual shifts. First, we compared the perceptual shifts for the conditions of view rubber hand at congruent posture versus view rubber hand at incongruent posture. The difference was significant, $t(7) = 2.9$, $p < .01$, one-tailed, suggesting that for the RHI to occur, the rubber hand had to be in a congruent position with respect to the participant's hand. Second, we compared the perceptual shifts for view rubber hand at congruent posture versus the condition of view wooden stick. The difference was significant, $t(7) = 4.311$, $p < .003$, one-tailed, suggesting that simple association of synchronous visuotactile events between the participant's hand and the neutral viewed object did not suffice to induce large proprioceptive drifts. In fact, the perceptual shifts for the view-wooden-stick conditions were not significantly different from zero, $t(7) = -2.04$, $p > .05$.

Overall, the results of Experiment 1 suggest that the RHI occurs only when the viewed rubber hand is in congruent posture with the participant's unseen hand. Moreover, the mere correlation of tactile and visual stimulation between one's own hand and a neutral object, such as a wooden stick, does not suffice for eliciting the RHI.

Experiment 2

Experimental Design

The aim of Experiment 2 was to investigate the relative importance of the two components of the RHI, namely, the influence of a general body scheme representation and the process of visual–tactile integration. To investigate the general top-down effect of body representation on the RHI, we manipulated the handedness identity of the viewed rubber hand. Would the RHI be present if participants were stimulated on their left hand while they were watching a right rubber hand being stimulated synchronously and at the same location? Moreover, to investigate the local character of the multisensory integration, we measured the proprioceptive drift for both the stimulated finger and for another unstimulated digit.

Participants were always stimulated on their left middle finger, and they saw a left or a right rubber hand being stimulated on its middle finger. The experimental factors were (a) the mode of stroking (synchronous vs. asynchronous), (b) whether participants saw a left or a right rubber hand (congruent vs. incongruent hand identity), and (c) the finger judged (middle vs. thumb). Because the relative position of the middle finger is the same for both left and right hands, we stroked the knuckle of the middle finger across conditions. The position of the thumb is symmetrically opposite between left and right hands. Therefore, judgments of the perceived position of the thumb might be especially sensitive to whether participants viewed a left or a right rubber hand, in comparison to judgments of middle finger position.

Method and Participants

There were, thus, four experimental conditions and four control conditions. In the experimental conditions, both hands were stroked synchronously, whereas in the control conditions, the two hands were stroked asynchronously. The conditions were blocked. Each participant performed the blocks in a different random order.

The experimental design and the setup are shown in Figure 3. Stimulation lasted 4 min. After the stimulation period, the rubber hand was covered, and a ruler was presented on a horizontal surface 18 cm above the participant's fingertips and aligned with the participant's frontoparallel plane. Participants were asked, “Where is your middle finger?” or “Where is your thumb?” and in response, they verbally reported a number on the ruler. Participants were trained to judge the position of their finger by projecting a parasagittal line from the center of their fingertip to the ruler. They repeated the judgment four times, in 15-s intervals. Other aspects of the method were the same as in Experiment 1.

Eight volunteers (mean age = 24 years, 8 female), with normal or corrected-to-normal vision, participated in this study after giving their informed consent.

Results

Figure 4 shows the mean proprioceptive drifts across conditions, averaged across the four posttest trials. Our analysis showed that the RHI was stationary over that period of 1 min. We found only very minor differences across these four estimates, particularly in the synchronous conditions. The range across these four estimates was 8 mm in the worst case and only 2 mm in the best case.

We submitted the proprioceptive drifts across conditions in a $2 \times 2 \times 2$ ANOVA with three within-subjects factors. The factors were (a) the rubber hand identity (congruent vs. incongruent), (b) the finger judged (middle vs. thumb), and (c) the mode of stroking (synchronous vs. asynchronous). There was a significant main effect of the finger judged, $F(1, 7) = 6.48$, $p < .05$, and of the mode of stroking, $F(1, 7) = 6.45$, $p < .05$. Moreover, the interaction of the three factors was significant, $F(1, 7) = 10.4$, $p < .05$.

We used simple effects analysis to investigate this pattern of interaction in more detail. This showed that differences between synchronous and asynchronous conditions were significant only for judgments of the middle finger position, when participants were looking at a congruent rubber hand, $t(7) = 3.89$, $p < .04$, two-tailed. Differences between synchronous and asynchronous stimulation for all the other conditions were not significant, $t(7) = 0.736$, $p > .05$, when judging thumb after looking at a congruent rubber hand, $t(7) = -0.279$, $p > .05$, when judging thumb after looking at an incongruent rubber hand, and $t(7) = -0.159$, $p > .05$, when judging middle after looking at an incongruent rubber hand. Therefore, significant proprioceptive drifts occurred only when participants watched a congruent rubber hand (i.e., left) and judged the position of the stimulated finger (i.e., middle).

To isolate the part of the positional drift due to visual–tactile integration and to obtain a true measure of RHI, we calculated the mean perceptual shifts as in Experiment 1. We subtracted the proprioceptive drifts obtained in the asynchronous conditions from the proprioceptive drifts obtained in the synchronous conditions. We investigated the resulting perceptual shifts using two planned comparisons.

We investigated whether attribution occurs only when seen and

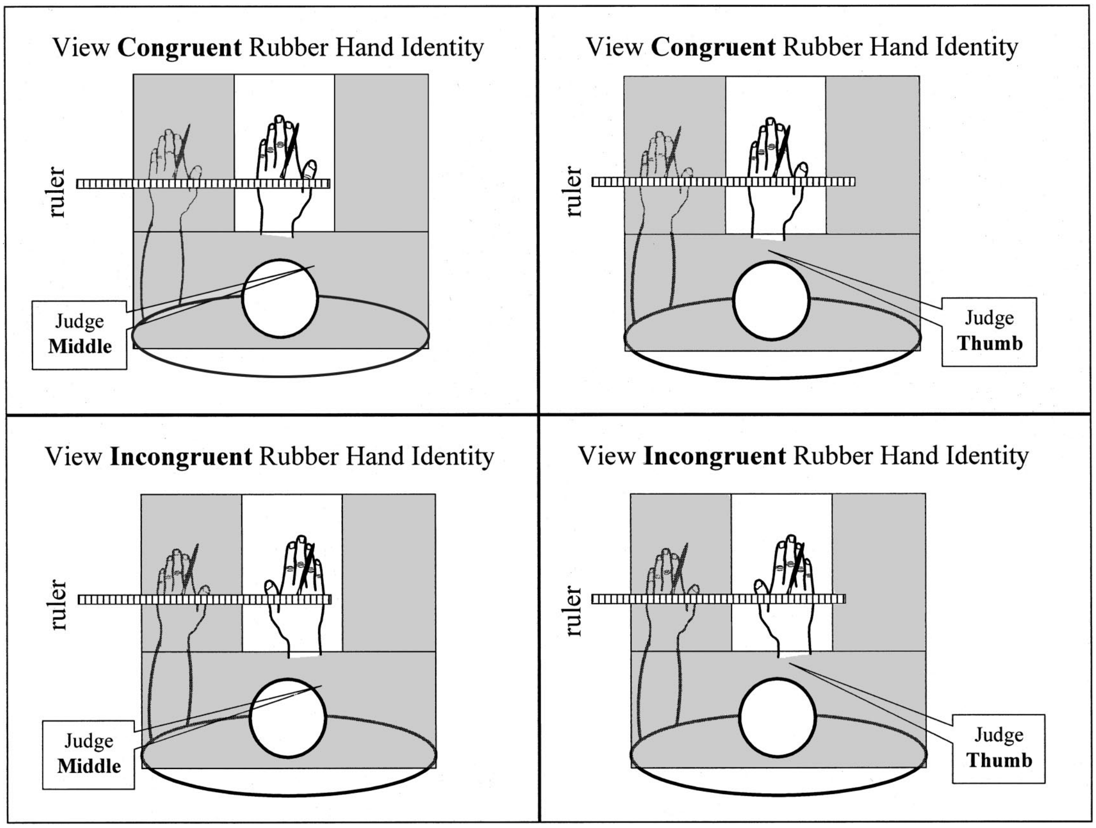

Figure 3. Experimental design and setup in Experiment 2.

felt hand identities are congruent. We compared the mean perceptual shifts for the middle finger when participants watched a left or a right rubber hand. The difference was significant, $t(7) = 2.9, p < .03$, two-tailed, suggesting that the RHI occurred only when felt and seen identities match were congruent. Because an incongruent rubber hand identity did not induce the RHI, our original prediction regarding the sensitivity of the thumb position was not confirmed. Judgments regarding the thumb position showed only minimal effect.

Nevertheless, we used a second planned comparison to investigate whether the proprioceptive drift was generalized across the whole hand or whether it was restricted to the locus of stimulation. We compared the perceptual shifts for the stimulated middle finger versus those for the unstimulated thumb when participants were looking at a congruent rubber hand. The difference was significant, $t(7) = 3.32, p < .02$, two-tailed, suggesting that only the stimulated finger was perceived to be closer to the rubber hand.

Overall, the results of Experiment 2 replicated the findings of Experiment 1 regarding the influence of general body-scheme representations on the RHI, because incongruent rubber hand identity, as was the case with the incongruent posture (Experiment 1), did not elicit the illusion. Furthermore, the analysis of Experiment 2 suggested that there is a localized process underlying the proprioceptive drift, because large shifts were observed only for the stimulated finger.

Experiment 3

Experimental Design

Experiment 3 was designed to further investigate the localized proprioceptive drifts of RHI seen in Experiment 2. The experimental design was a $2 \times 2 \times 2$ factorial. The three factors were (a) the mode of stroking (synchronous vs. asynchronous), (b) the finger stroked (index vs. little), and (c) the finger judged (index vs. little). Only one finger was stroked per condition, always the same as the one on the rubber hand. Participants were asked to judge the position of either the stimulated or the unstimulated finger. Only the fingers of the left hand were stimulated, and participants always watched a left rubber hand, being stroked at the exact same location as their own left hand.

Method and Participants

There were four experimental (i.e., synchronous) and four control (i.e., asynchronous) conditions. The conditions were blocked and randomly ordered between participants. Stimulation was delivered horizontally on the index finger or the little finger. Other aspects of the methods were the same as in Experiment 2.

Ten right-handed volunteers (mean age = 23 years, 7 female and 3 male), with normal or corrected-to-normal vision, participated.

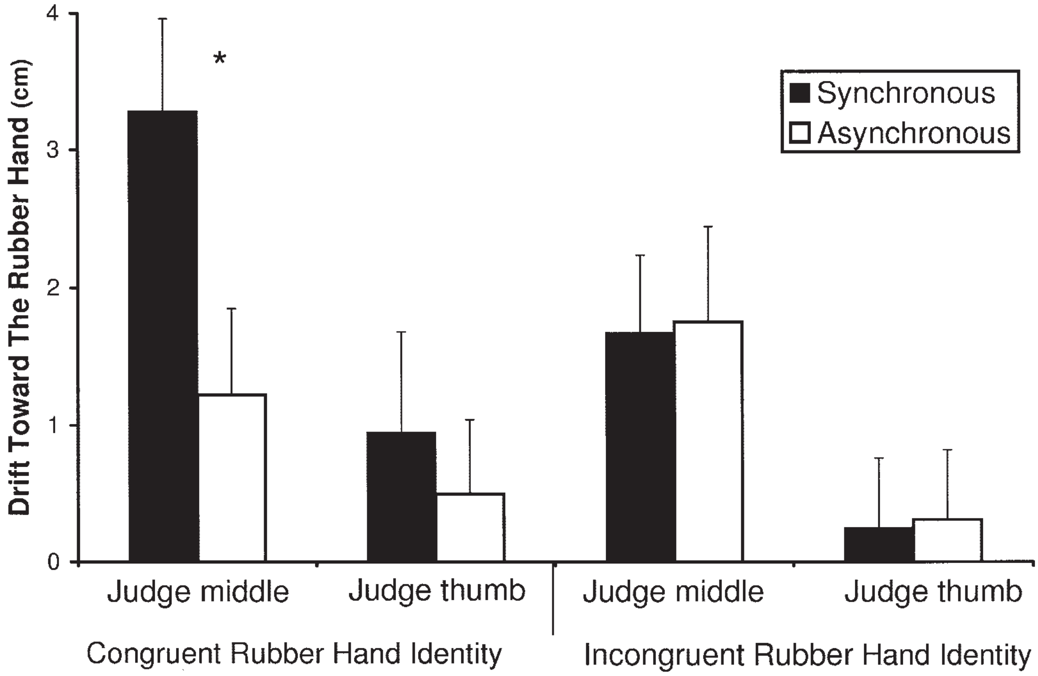

Figure 4. Mean proprioceptive drifts toward the rubber hand in Experiment 2. Error bars indicate standard errors. Asterisk indicates significant differences between synchronous and asynchronous stimulation.

Results

Simple effects analysis was used to compare the mean proprioceptive drifts between the synchronous and the asynchronous conditions. Differences were significant only when participants judged the finger that had just been stroked: for the index finger, $t(9) = 3.983$, $p < .05$, and for the little finger, $t(9) = 4.082$, $p < .05$. When the finger judged was not the finger stroked, the drifts were not significantly different for the two different modes of stroking, $t(9) = 0.637$, $p > .05$, for the index finger, and $t(9) = 0.746$, $p > .05$, for the little finger, suggesting that the RHI did not occur.

By subtracting the proprioceptive drifts obtained for the asynchronous conditions from the synchronous ones, we again calculated the perceptual shifts. Figure 5 shows the mean perceptual shifts across conditions. Only the perceptual shifts for the conditions in which the finger judged was the finger stroked were reliably greater than zero, $t(9) = 4.082$, $p < .002$, one-tailed, for the index finger, and $t(9) = 3.98$, $p < .002$, one-tailed, for the little finger.

The mean perceptual shifts were also submitted to a 2 × 2 repeated measures ANOVA with two within-subject factors. The first factor was the finger stroked (index or little), and the second factor was the finger judged (index or little). The main effect of the finger stroked was not significant, $F(1, 9) = 0.047$, $p > .05$, nor was the main effect of the finger judged, $F(1, 9) = 0.118$, $p > .05$. However, the interaction of the two factors was significant, $F(1, 9) = 6.074$, $p < .05$, because larger shifts occurred when judging the position of the stimulated finger compared with the unstimulated finger.

Experiment 4

Experimental Design

The pattern of localized perceptual shifts that we observed in Experiment 3 could in principle arise because only one finger was stimulated in each condition. We therefore performed Experiment 4, which provided a better control for possible attentional artifacts. To achieve this, we always stimulated two fingers on each trial, but we manipulated the pattern of stimulation on each finger independently. We also used computer-controlled motors to ensure that the overall amount of stimulation was precisely matched across conditions.

In all conditions, we stroked the index and the little fingers of the participant's left hand and of a left rubber hand. The conditions differed only in the pattern of stimulation across the fingers (i.e., the index finger and the little finger). The rubber hand always received the same pattern of stimulation, to the index and little fingers simultaneously. However, the participant's index or little finger might or might not receive the same pattern of stimulation as the rubber hand. Thus, there were four different stimulation conditions: (a) both index and little fingers were stroked in synchrony with the index and little fingers of the rubber hand, (b) only the participant's index finger was stroked in synchrony with the rubber fingers, while the little finger was stroked asynchronously, (c) only the participant's little finger was stroked in synchrony with the rubber fingers, while the index was stroked asynchronously, and (d) both the participant's fingers were stroked asynchronously with respect to the rubber fingers. For each stimulation condition, participant had to judge, in different blocks, the position of (a) their index, (b) their little, or (c) their middle finger. Thus,

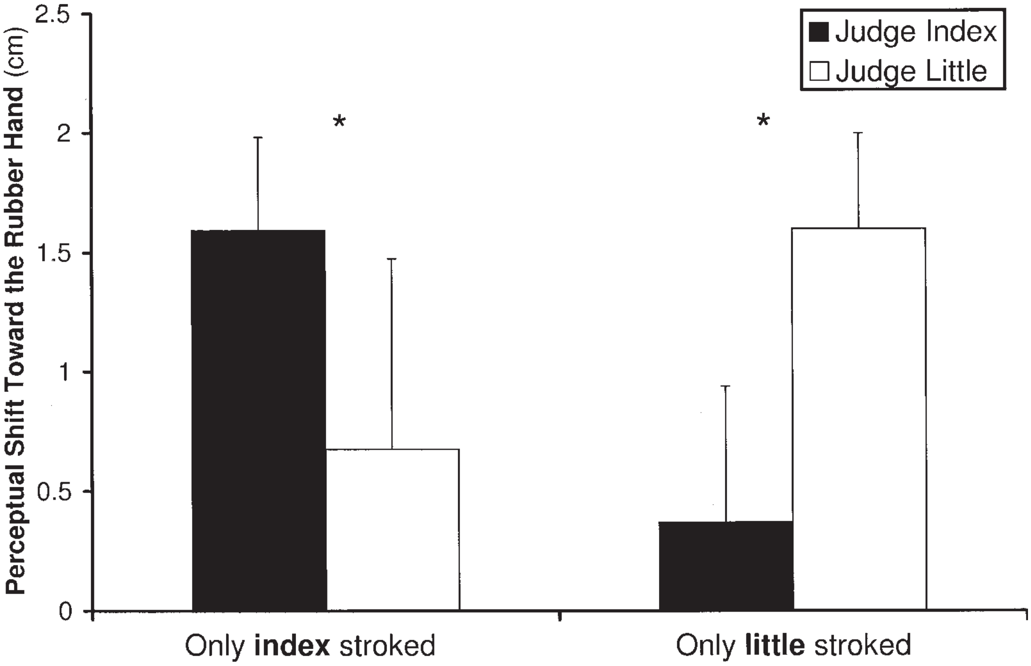

Figure 5. Mean perceptual shifts in Experiment 3. Error bars indicate standard errors. Asterisks indicate significant differences between judgments for the stimulated and the unstimulated finger.

the experimental factors were the mode of stroking and the finger judged. The experiment consisted of 12 conditions that were blocked and randomly ordered between participants.

Method and Participants

Stimulation was delivered mechanically by two stepper motors. To simulate the unpredictable nature of manual stroking, we randomly varied the direction and speed of stroking within conditions. The participant's hand and the rubber hand always received the same random direction and speed on each stroke. In the synchronous condition, the participant's hand and the rubber hand were stroked simultaneously. In the asynchronous conditions, the finger was stroked with a random delay of 500–1,000 ms after stroking of the rubber hand. The total amount and spatial pattern of stimulation was the same across all conditions. Synchronous and asynchronous conditions differed only in the degree of temporal correlation of visual and tactile stimulation.

Figure 6 shows the experimental setup. Participants sat at a table. The rubber hand was presented in front of them, aligned with their midline. Participants could see the rubber hand through a two-way mirror. Before the start of each block, the rubber hand was invisible and participants made one pretest baseline judgment for the position of the indicated finger. Stroking lasted 5 min per condition. Participants made one judgment per minute. During the judgments, there was no tactile stimulation, and the lights under the two-way mirror were switched off to make the rubber hand invisible. Other aspects of the method were the same as in Experiment 1.

Fourteen volunteers (mean age = 26.5 years, 11 female and 3 male), with normal or corrected-to-normal vision and all right-handed, gave their informed consent to participate in this study.

Results

The mean raw judgment errors are shown in Figure 7 as a function of time. Across all conditions, a positive judgment error was found even before stroking began. That is, participants perceived their hand to be closer to the rubber hand than it truly was, even without stimulation, as in Experiment 1. For the synchronous conditions of stroking, the drifts toward the rubber hand were rapid, especially during the first 2 min. This suggests that the main effect of visuotactile integration was on the early portion of the curve. In asynchronous conditions, this early rise was reduced or absent. We averaged the mean judgment errors for the five judg-

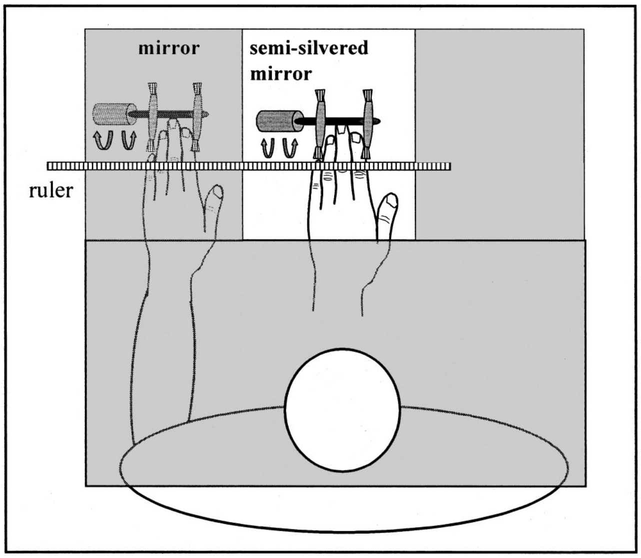

Figure 6. Experimental setup in Experiment 4.

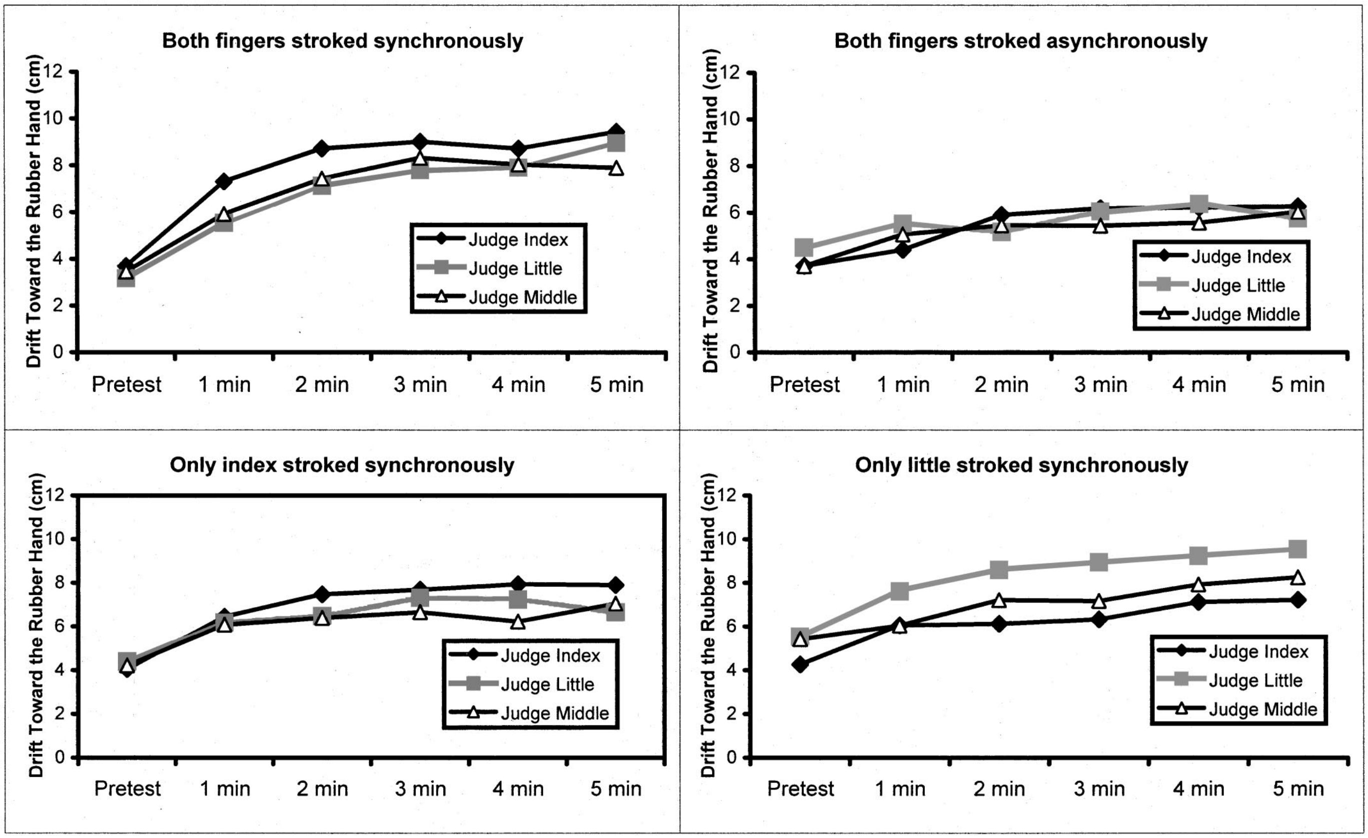

Figure 7. Mean raw judgment errors in the perceived position of each finger across time in Experiment 4. Participants made one pretest judgment and then four judgments every 1 min. Zero represents the real position of the participant's hand.

ments obtained during and after the stroking period and subtracted the baseline pretest judgment. We subjected the resulting proprioceptive drifts were to two separate analyses.

The first analysis investigated the local effect of attribution. To investigate the local effect of visual–tactile correlation on the proprioceptive drift, we performed a 2 × 2 ANOVA on the data of Figure 8 with two within-subjects factors. The factors were the mode of stroking (only the index finger synchronous vs. only the little finger synchronous) and the finger judged (index vs. little). Neither the main effect of mode of stroking, $F(1, 13) = 0.144$, $p >$ .05), nor the main effect of finger judged, $F(1, 13) = 0.004$, $p > .05$, was significant. But, the interaction between them was significant, $F(1, 13) = 5.26$, $p < .05$. We observed this interaction because drifts toward the rubber hand were greatest when the finger judged was the one that was stroked synchronously with the rubber finger. Therefore, this crossover pattern of results replicates Experiment 3, with the additional control that all fingers received equal amounts of stimulation in all conditions.

Our second analysis focused on the global effect of attribution. We performed a 2 × 3 ANOVA on the mean judgment errors of Figure 9, with two within-subjects factors. The factors were the mode of stroking (both fingers synchronously vs. both fingers asynchronously) and the finger judged (index, little, or middle). Only the main effect of mode of stroking was significant, $F(1, 13) = 15.5$, $p < .01$, whereas the main effect of finger judged was not significant, $F(2, 26) = 1.19$, $p > .05$. The interaction of the two factors did not reach significance, $F(2, 26) = 0.832$, $p > .05$. When both fingers were stimulated at the same time as the rubber fingers, the unstimulated middle finger drifted toward the rubber hand. The drift for the unstimulated middle finger was similar to those for the stimulated index and little fingers. A planned comparison for the judged position of the middle finger in the synchronous versus asynchronous conditions confirmed that the RHI occurred in the absence of stimulation, $t(13) = 3.32$, $p < .01$, two-tailed.

Taken together, these two analyses show that the RHI involves partly a bottom-up process (because the shift is greater for a synchronously stimulated finger than for an asynchronously stimulated one), and partly a top-down process to preserve an overall coherence of nonvisual body representation (because an unstimulated finger showed the drift in the perceived position when its neighbors were stimulated synchronously with respect to the rubber fingers).

General Discussion

The RHI was originally reported by Botvinick and Cohen (1998). When participants viewed a rubber hand being stimulated

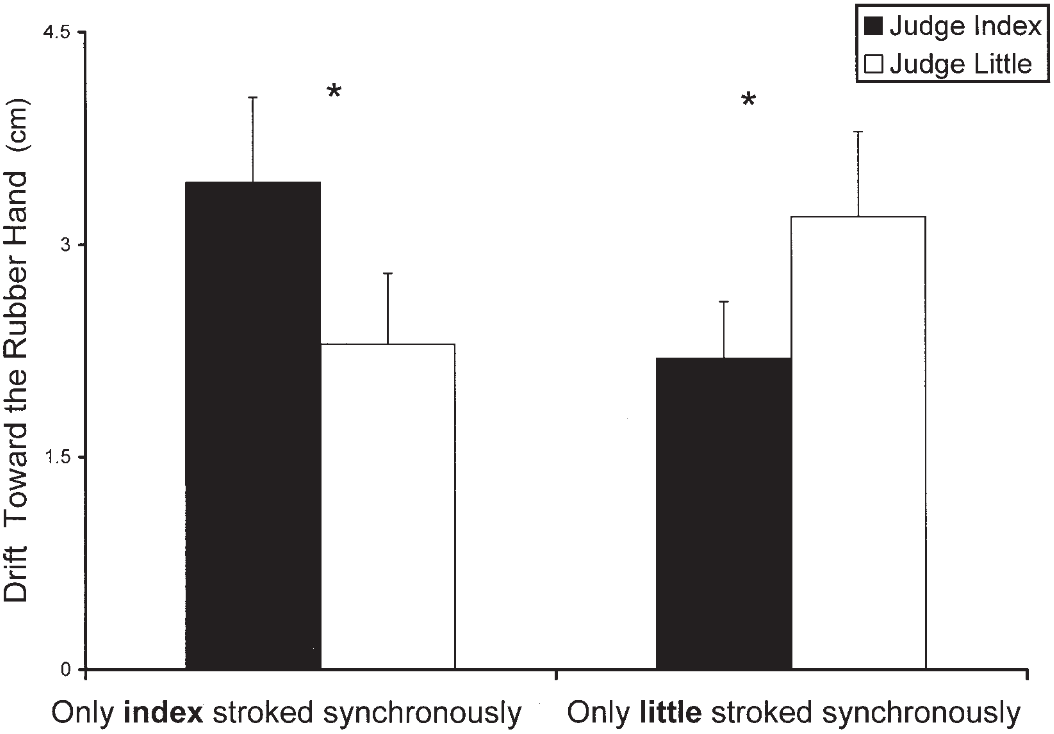

Figure 8. Mean proprioceptive drifts toward the rubber hand for the index and little fingers in Experiment 4. Error bars indicate standard errors. Asterisks indicate significant differences between judgments for the synchronously stimulated finger and the asynchronously stimulated finger. The crossover pattern of results replicates that of Experiment 3 (see Figure 5).

in synchrony with their own unseen hand, they misperceived the position of their hand as being closer to the rubber hand than it really was. Our results replicate the illusion originally reported by Botvinick and Cohen, and they are compatible with previous research on crossmodal integration in the RHI (Farné, Pavani, Meneghello, & Ladavas, 2000; Rorden, Heutnik, Greenfield, & Robertson, 1999).

We reported a series of four experiments on the RHI. In Experiments 1 and 2, we focused on the role of body scheme representations on the RHI. Our results suggested that the RHI occurred only when the rubber hand was in a congruent posture or of a congruent identity with respect to the participant's hand. Incongruent rubber hand posture, incongruent rubber hand identity, and neutral objects did not elicit similar proprioceptive drifts. These findings implied that mere correlation between visual and tactile percepts is not a sufficient condition for self-attribution of a rubber hand. Another significant result was the pattern of proprioceptive drifts. In Experiment 3, only the felt position of the stimulated finger drifted toward the rubber hand. In Experiment 4, we showed that even when two fingers received an equal amount of stimulation, it was only the felt position of a synchronously stimulated finger that drifted significantly toward the rubber hand. These findings showed that synchronous visual and tactile correlation is a necessary condition for the inducement of the RHI. However, this pattern of localized proprioceptive drifts was not absolute. In Experiment 4, the unstimulated middle finger drifted by an amount similar to the drift in other fingers, which were stimulated. This finding suggests that there is a spreading gradient of the RHI to unstimulated fingers. Before discussing the implications of our results and their relation to previous studies, we consider certain methodological issues.

There are several methodological differences between the original experiment by Botvinick and Cohen (1998) and our experiments. First, in the Botvinick and Cohen experiment, synchronous or asynchronous stimulation was a between-subjects factor, whereas in the present study we manipulated it as a within-subjects factor. Asynchronous stimulation served as baseline, providing a better control for the effects of physical stimulation. Second, Botvinick and Cohen asked the participants to make intermanual reaches with their unstimulated hand to the felt position of their index finger. In our experiments, participants judged the position of their fingers by indicating a number on a ruler presented in front of them, and both their hands were kept still during judgment. In our study, the rubber hand was aligned with the participant's midline, whereas in the original experiment, the rubber hand was positioned next to the participant's stimulated left hand. Botvinick and Cohen stimulated the hands using a random pattern of movements "distal to the wrist, [on] the dorsal aspect of hand and fingers, including thumb, and fingertips" (M. Botvinick, personal communication, October 16, 2002). We stroked the hand only along the fingers, from the knuckle to the fingertip. We stimulated only one finger per condition in Experiments 1, 2, and 3, and two fingers in Experiment 4. This pattern enabled us not only to control

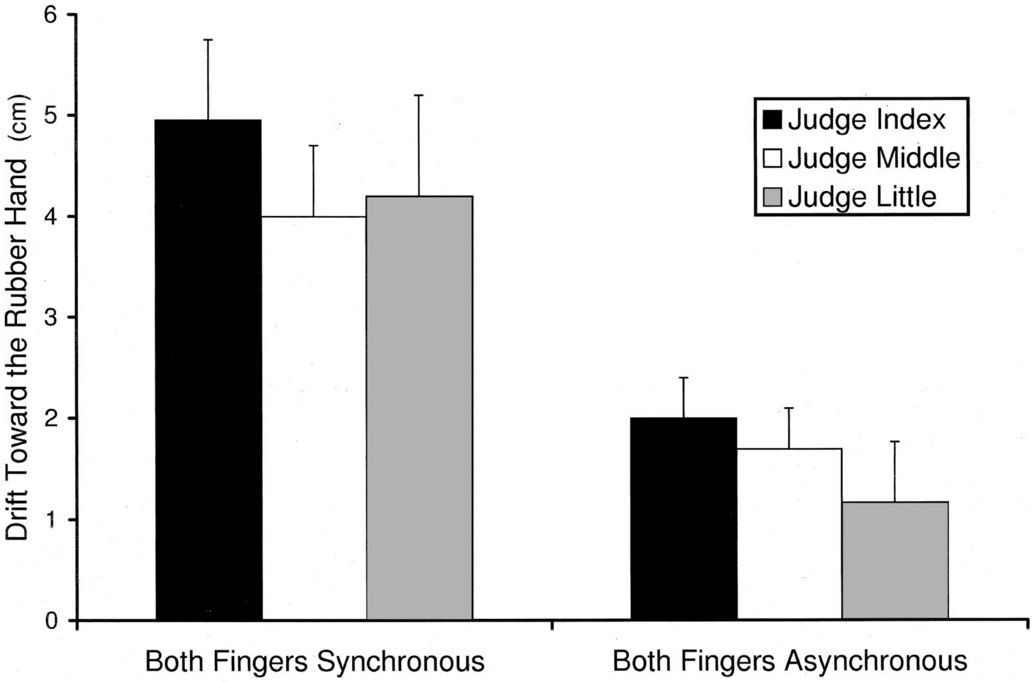

Figure 9. Mean proprioceptive drifts toward the rubber hand for the index, little, and middle fingers in Experiment 4, in which both index and little fingers were stimulated synchronously or asynchronously. Error bars indicate standard errors.

the amount of stimulation across conditions but also to investigate the localized multisensory integration.

What is the mechanism underlying the self-attribution of the rubber hand to one's own body? Botvinick and Cohen (1998) suggested that intermodal matching is a sufficient condition for "self-attribution" of the rubber hand to one's own self. Other studies (e.g. Armel & Ramachandran, 2003) reported that it is possible to induce a sense of ownership by stroking the participants' unseen hand while they viewed a neutral object (such as a shoe or a table) being synchronously stroked. Armel and Ramachandran (2003) concluded that this illusion ("it feels like the fake hand/table is my hand," p. 1504) is the result of Bayesian perceptual learning. That would suggest that the RHI is the result of a purely bottom-up mechanism, which associates synchronous visuotactile events. In the strong version of this model, any object can become part of "me," simply because strong statistical correlations between different sensory modalities are both necessary and sufficient conditions for self-attribution. On this view, psychological concepts such as embodiment and selfhood are unnecessary because purely Bayesian principles of statistical correlation are sufficient to extend the body representation to include even body parts as implausible as tables (Armel & Ramachandran, 2003). In fact, Armel and Ramachandran argued that the RHI is resistant to top-down knowledge such as cognitive body representations.

This Bayesian account of the RHI predicts that (a) the illusion should be present even when participants are looking at a neutral object that has no functional relationship with the stimulated body part, (b) the illusion should be present even when the rubber hand is of incongruent posture or identity, (c) proprioceptive drifts should be significantly larger only for the stimulated digits because strong statistical correlations will be established only for the stimulated fingers, and (d) proprioceptive drifts should be smaller or even absent for unstimulated body parts because no statistical correlation between visual and tactile experience exists in that case. The results of Experiments 1–4 that are reported in this article challenge this model. First, it seems that attribution requires a plausible and congruent visual object (i.e., a rubber hand) to bind with a body part, with respect to the general body configuration (i.e., posture, hand identity). Second, the pattern of proprioceptive drifts reported in this article challenges the Bayesian model of the RHI. We examine each point separately.

With regard to the role of general body gestalts, the RHI occurred only when participants viewed a congruent rubber hand that was stimulated synchronously with their own hand. In these conditions, participants perceived their hand to be closer to the rubber hand than it really was. When the rubber hand was replaced by a neutral object (Experiment 1), the RHI did not occur. Therefore, synchronous visual and tactile stimulation is not a sufficient condition for the inducement of the RHI, and the suggestion that the illusion is the product of a purely bottom-up process should be ruled out. Moreover, when the rubber hand was positioned in an incongruent posture (Experiment 1), or when it was of an incongruent identity (Experiment 2), the RHI did not occur. These findings suggest that even when participants were looking at a rubber hand, the synchronicity of visual and tactile events did not suffice for the inducement of the RHI because of the incongruence between the rubber hand and the participants' own hand at a level that goes beyond synchronous stimulation (Experiments 1 and 2). Hand posture and hand identity have been identified as two kinds of body representations that modulate the visuotactile integration underlying the RHI. Indeed, we found an inverse RHI effect when participants viewed incongruent objects (Experiment 1). In these conditions, the perceptual shifts were in an opposite direction, namely, away from the rubber hand. This "perceptual repulsion" might reflect a behavioral correlate of the dissociation between events that either do or do not produce self-attribution. Interestingly, studies of temporal awareness of action have also described a perceptual attraction between one's own action and an external consequence, and a perceptual repulsion for a physically comparable involuntary movement followed by the identical external consequence (for a review, see Tsakiris & Haggard, in press). In both body representation and intentional action, perceptual attraction seems to be the hallmark of the bodily self, whereas perceptual repulsion may be characteristic of events not linked to the self. Overall, the results of Experiments 1 and 2 suggest that visual body representations can best be combined with the participant's preexisting body scheme when both temporal (Graziano, Cooke, & Taylor, 2000) and spatial (van den Bos & Jeannerod, 2002) cues are congruent. Contrary to the Bayesian model of RHI, according to which concurrent visual and tactile simulation *constructs* a changed body scheme (Armel & Ramachandran, 2003), we suggest that the concurrent visuotactile inputs are integrated within a preexisting representation of one's own body. Neurophysiological studies on monkeys support this view. Graziano and colleagues (Graziano et al., 2000) showed that bimodal neurons in parietal Area 5 of the monkey brain were sensitive to the position of a fake arm when fake and real hands were stroked synchronously, but only when the fake arm was aligned with the monkey's body. Moreover, these bimodal neurons were not sensitive when the fake arm was of an incongruent identity, or placed at an incongruent posture, or when a neutral object replaced the fake hand.

It might be argued that the pattern of localized proprioceptive drifts provides support for the Bayesian approach because stronger statistical correlations are expected only for the stimulated finger or fingers. The strong version of the Bayesian model postulates that the body scheme is constructed on the basis of visuotactile integration. Therefore, the pattern of localized proprioceptive drifts, which reflects a distortion of the hand scheme, can be explained by the nature of localized visuotactile stimulation used in our experiments. Nevertheless, this strong trend for localized drifts was restricted in two cases: First, the stimulated middle finger did not drift when participants were looking at an incongruent rubber hand identity (Experiment 2), and second, the unstimulated middle finger drifted to an equivalent extent when two fingers were stimulated synchronously (Experiment 4). In Experiment 4, the unstimulated middle finger drifted more toward the rubber hand when the neighboring fingers were stroked synchronously than when they were stroked asynchronously. A purely bottom-up account cannot explain how this effect was possible in the absence of stimulation. There seems to be a gradient of generalization of the illusion to the adjacent unstimulated fingers. This implies an influence of an additional nonvisual representation of body parts, which mediates the altered proprioception in the RHI. We suggest that the localized proprioceptive drifts result from bottom-up processes, whereas the spreading pattern observed in Experiment 4 results from top-down influences. This interpretation is compatible with a series of neurophysiological studies showing that finger representations in primary somatosensory cortex are dynamically modulated by both top-down and bottom-up influences after prolonged durations of stimulation. Depending on the synchronous or asynchronous stimulation of multiple fingers, cortical representations of fingers can be either integrated or segregated (Braun, Schweizer, Elbert, Birbaumer, & Taub, 2000, Braun, Wilms, et al., 2000, Braun et al., 2002; Wang, Merzenich, Sameshima, & Jenkins, 1995). Braun and colleagues demonstrated that when participants were passively attending to the synchronous stimulation of digits one and five, the distance between the cortical representations of the two digits decreased as measured by neuroelectric source localization of electroencephalograph. However, when participants were asked to discriminate stimulation on these fingers, the distance between the cortical representations of the same digits was increased compared with the passive stimulation conditions. This finding suggests that cortical representations of the fingers are also modulated by top-down factors, such as the discrimination task used in that study.

In relation to the phenomenological attribution of the rubber hand to the bodily self, it is still unclear what kind of objects can be "incorporated" into one's own body. Research on tool use has demonstrated the neurophysiological mechanisms underlying incorporation of tools in both primates (Iriki et al., 1996) and humans (Berti & Frassinetti, 2000; Farné & Ladavas, 2000; Yamamoto & Kitazawa, 2001a, 2001b). Given the plasticity of the somatosensory cortex and the functional role of bimodal neurons in parietal Area 5, it is possible that one's own body scheme can be extended so as to include even one's own car (for a review, see Graziano & Botvinick, 2002)! It has been suggested that the first stage at which visual information about arm position is integrated with somatosensory information seems to be in Area 5 (Graziano, 1999; Graziano et al., 2000). Therefore, the functional role of the bimodal neurons found in Area 5 would be that of "a central node in representing the configuration of the body" (Graziano & Botvinick, 2001, p. 143). It is expected that these bimodal neurons in the human homologue area will have a specific role in underpinning the inducement of the RHI.

Extensibility of the body scheme involves an interaction between bottom-up and top-down processes. First, we showed that the correlation of sensory inputs in time (i.e., synchronicity) is the main driver of the RHI (all experiments). Second, we showed that extensibility and adaptation of body representation depend on coherence with preexisting visual (Experiments 1 and 2) and proprioceptive (Experiments 1 and 4) body representations. Our results favor an interplay between bottom-up and top-down influences in the process of bodily synthesis and self-attribution. We showed that intermodal matching is not a sufficient condition for self-attribution of a rubber hand. What, then, is the relationship between intermodal matching and the self? Developmental studies suggest that intermodal matching seems to be a prerequisite for self-identification (Rochat & Striano, 2000). However, the present results suggest that the phenomenal content associated with the perception of the rubber hand is more than a mere sum of correlated visual and tactile inputs. Sensory inputs related to the body seem to be integrated against a set of conditions that guarantee the functional and phenomenological coherence of bodily experience. In this sense, the cognitive representation of the body reflects the interplay between both sensory input and conceptual interpretation. In favor of this view, several studies on cross-modal effects have reported that viewing one's body via a video image or a mirror is sufficient to induce a range of visual–tactile interactions (Maravita, Spence, Sergent, & Driver, 2002; Tipper et al., 2001). That is, the functional relation of the viewed image to the tactile stimulus is important for cross-modal integration, but the precise visual properties may not be.

Experiments on the RHI can address empirically the relation to one's own body at different levels of functional description, from neural to phenomenological. The results presented in this article could also have important technological applications, for example, in the field of telepresence and sensory–motor experiences in virtual reality environments. Recent research on telepresence and teleoperation has addressed the role of the "sense of presence" in such virtual environments (Welch, 1999). The interplay between bottom-up and top-down influences in the generation of experiences such as the one found in the RHI demonstrates the need for a reconsideration of the cognitive representation of the body and of the ways that the living body can be "incorporated" in a virtual environment.

To conclude, the synchronized visual and tactile stimulation causes the illusion, but the phenomenological content within the illusion is a description of one's own body and not a description of the stimulation. Across four experiments, we showed that the spatiotemporal pattern of stimulation influences the RHI, that the induced shift is not an adaptation of limb position but a mislocalization of the stimulated body part, and that this mislocalization is constrained by an abstract representation of a coherent body scheme. Taken together, the results suggest that at the level of the process underlying the build up of the illusion, bottom-up processes of visuotactile correlation drive the illusion as a necessary, but not sufficient, condition. Conversely, at the level of the phenomenological content, the illusion is modulated by top-down influences originating from the representation of one's own body.

References

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

Berti, A., & Frassinetti, F. (2000). When far becomes near: Remapping of space by tool use. *Journal of Cognitive Neuroscience, 12,* 415–420.

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

Braun, C., Jaug, M., Wiech, K., Birbmaumer, N., Elbert, T., & Roberts, L. E. (2002). Functional organization of primary somatosensory cortex depends on the focus of attention. *NeuroImage, 17,* 1451–1458.

Braun, C., Schweizer, R., Elbert, T., Birbaumer, N., & Taub, E. (2000). Differential activation in somatosensory cortex for different discrimination task. *Journal of Neuroscience, 20,* 446–450.

Braun, C., Wilms, A., Schweizer, R., Godde, B., Preissl, H., & Birbaumer, N. (2000). Activity patterns of human somatosensory cortex adapt dynamically to stimulus properties. *NeuroReport, 11,* 2977–2980.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433.

Farné, A., & Ladavas, E. (2000). Dynamic size-change of hand peripersonal space following tool use. *NeuroReport, 11,* 1645–1649.

Farné, A., Pavani, F., Meneghello, F., & Ladavas, E. (2000). Left tactile extinction following visual stimulation of a rubber hand. *Brain, 123,* 2350–2360.

Ghilardi, M. F., Gordon, J., & Ghez, C. (1995). Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. *Journal of Neurophysiology, 73,* 2535–2539.

Graziano, M. S. (1999). 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.

Graziano, M. S., & Botvinick, M. M., (2002). How the brain represents the body: Insights from neurophysiology and psychology. In W. Prinz & B. Hommel (Eds.), *Common mechanisms in perception and action: Attention and performance XIX* (pp. 136–157). Oxford, England: Oxford University Press.

Graziano, M. S., Cooke, D. F., & Taylor, C. S. (2000). Coding the location of the arm by sight. *Science, 290,* 1782–1786.

Haggard, P., Taylor-Clarke, M., & Kennett, S. (2003). Tactile perception, cortical representation and the bodily self. *Current Biology, 13,* R170–R173.

Howell, D. C. (1997). *Statistical methods for psychology* (4th ed.). Belmont, CA: Duxbury.

Iriki, A., Tanaka, M., & Iwamura, Y. (1996). Coding of modified body schema during tool use by macaque postcentral neurones. *NeuroReport, 7,* 2325–2330

Kennett, S., Taylor-Clarke, M., & Haggard, P. (2001). Noninformative vision improves the spatial resolution of touch in humans. *Current Biology, 11,* 1188–1191.

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

Maravita, A., Spence, C., Sergent, C., & Driver, J. (2002). Seeing your own touched hands in a mirror modulates cross-modal interactions. *Psychological Science, 13,* 350–355.

Merleau-Ponty, M. (1962). *Phenomenology of perception* (C. Smith, Trans.). London: Routledge. (Original work published 1945)

Pavani, F., Spence, C., & Driver, J. (2000). Visual capture of touch: Out-of-the-body experiences with rubber gloves. *Psychological Science, 11,* 353–359.

Ramachandran, V. S., & Hirtsein, W. (1998). The perception of phantom limbs. *Brain, 121,* 1603–1630.

Rochat, P., & Striano, T. (2000). Perceived self in infancy. *Infant Behavior & Development, 23,* 513–530.

Rorden, C., Heutnik, J., Greenfield, E., & Robertson, I. H. (1999). When a rubber hand "feels" what the real hand cannot. *NeuroReport, 10,* 135–138.

Tipper, S. P., Phillips, N., Dancer, C., Lloyd, D., Howard, L. A., & McGlone, F. (2001) Vision influences tactile perception at body sites that cannot be viewed directly. *Experimental Brain Research, 139,* 160–167.

Tsakiris, M., & Haggard, P. (in press). Experimenting with the acting self. *Cognitive Neuropsychology*.

van den Bos, E., & Jeannerod, M. (2002). Sense of body and sense of action both contribute to self-recognition. *Cognition, 85,* 177–187.

Wang, X., Merzenich, M. M., Sameshima, K., & Jenkins, W. M. (1995). Remodelling of hand representation in adult cortex determined by timing of tactile stimulation. *Nature, 378,* 71–75.

Welch, R. B. (1999). How can we determine if the sense of presence affects task performance? *Presence: Teleoperators and Virtual Environments, 8,* 574–577.

Welch, R. B., Widawski, M. H., Harrington, J., & Warren, D. H. (1979). An examination of the relationship between visual capture and prism adaptation. *Perception & Psychophysics, 25,* 126–132.

Yamamoto, S., & Kitazawa, S. (2001a). Reversal of subjective temporal order due to arm crossing. *Nature Neuroscience, 4,* 759–765.

Yamamoto, S. & Kitazawa, S. (2001b). Sensation at the tips of invisible tools. *Nature Neuroscience, 4,* 979–980.

Received November 26, 2003
Revision received June 21, 2004
Accepted June 24, 2004 ■