Grasping Visual Illusions: No Evidence for a Dissociation Between Perception and Action

V.H. Franz, K.R. Gegenfurtner, H.H. Bülthoff and M. Fahle
Psychological Science 2000 11: 20
DOI: 10.1111/1467-9280.00209

Research Article

GRASPING VISUAL ILLUSIONS: No Evidence for a Dissociation Between Perception and Action

V.H. Franz,[1] K.R. Gegenfurtner,[1] H.H. Bülthoff,[1] and M. Fahle[2]

[1]Max-Planck-Institut für Biologische Kybernetik, Tübingen, Germany, and [2]Zentrum für Kognitionswissenschaften, Universität Bremen, Bremen, Germany

Abstract—*Neuropsychological studies prompted the theory that the primate visual system might be organized into two parallel pathways, one for conscious perception and one for guiding action. Supporting evidence in healthy subjects seemed to come from a dissociation in visual illusions: In previous studies, the Ebbinghaus (or Titchener) illusion deceived perceptual judgments of size, but only marginally influenced the size estimates used in grasping. Contrary to those results, the findings from the present study show that there is no difference in the sizes of the perceptual and grasp illusions if the perceptual and grasping tasks are appropriately matched. We show that the differences found previously can be accounted for by a hitherto unknown, nonadditive effect in the illusion. We conclude that the illusion does not provide evidence for the existence of two distinct pathways for perception and action in the visual system.*

Several theories state that visual information is processed in two different streams in the primate brain—the dorsal and the ventral streams. Based on lesion studies on monkeys, Ungerleider and Mishkin (1982) proposed that the function of the dorsal stream is the analysis of the spatial relations between objects ("where" pathway), and the function of the ventral stream is the recognition of objects ("what" pathway).

On the basis of neuropsychological studies (Goodale, Milner, Jakobson, & Carey, 1991; Goodale et al., 1994) showing a double dissociation between grasping an object and perceiving its shape, Goodale and Milner (1992; Milner & Goodale, 1995) reinterpreted this theory. They suggested that the function of the dorsal stream is not to analyze the location of objects, but rather to guide the manipulation of objects, whereas the function of the ventral stream is to perform computations that are necessary for object recognition and conscious perception. They argued that the computations for these functions must fulfill totally different requirements. Computations for the guidance of an action have to be fast, have to precisely code the position of the object relative to the effector, and need only a short memory because the position of the object can change quickly. In contrast, computations for the purposes of object recognition do not need to be as precise and fast. In this case, it is much more important to evaluate an object in its context and to enable a long-lasting representation. Therefore, Milner and Goodale (1995) suggested that there exist two different visual systems, one that guides motor actions and one that leads to object recognition and conscious perception. Of course, it would be strong evidence if such a division of labor were reflected not only in deficits of neuropsychological patients, but also in the healthy visual system. Besides other psychophysical evidence (e.g., Bridgeman, Lewis, Heit, & Nagle, 1979; Goodale, Pélisson, & Prablanc, 1986; Hansen & Skavenski, 1985), especially strong support seemed to come from studies on grasping in humans.

To grasp, humans have to move their hand close to the target object. During the reach, the index finger and thumb open to a maximum aperture (Fig. 1a) that is linearly related to object size (Fig. 1b; Jeannerod, 1981, 1984). This maximum preshape aperture is formed before the hand has any contact with the object. Therefore, the maximum preshape aperture reflects the size estimate being transferred to the motor system from the visual system (if no other, nonvisual cues about object size are available). Because conscious perception is receptive to a number of size illusions (Coren & Girgus, 1978), the question arises whether the maximum preshape aperture will be affected by these illusions as well. In principle, there are two possibilities: The first is that the motor system uses the same internal representation of object size that perception does. This common representation would be influenced by the illusion. Because of the linear relationship between object size and maximum preshape aperture (Fig. 1b), it would be possible to predict the influence of the illusion on maximum preshape aperture. The second possibility is that the motor system uses a different representation of object size than perception does. In this case, the maximum preshape aperture could be unaffected by the illusion. This is the prediction of the perception-versus-action hypothesis of Milner and Goodale (1995). The computations for perception should focus on the relationship between an object and its surrounding objects, whereas the computations for action should focus on the relationship between an object and the effector to be used (in this case, the hand). Because visual size illusions are often generated by special arrangements of objects, the theory predicts that these illusions should have little (or no) influence on the motor system.

The first and most influential study investigating this topic was performed by Aglioti, DeSouza, and Goodale (1995) and was replicated by Haffenden and Goodale (1998). These original studies used the Ebbinghaus (or Titchener) illusion: A central circle appears smaller when surrounded by large circles than when surrounded by small circles. So that the influence of this illusion on grasping could be determined, the central circle was replaced with a disc, which was grasped by the subjects. Results showed a larger effect of the illusion on perception than on the maximum preshape aperture. These results were interpreted as strong evidence for the theory of Milner and Goodale (1995).

Several problems in the original studies prompted us to try a replication using an improved and simplified design: First, even though the effect on perception was larger, the original studies also found influences on grasping—as did other studies (Brenner & Smeets, 1996; Daprati & Gentilucci, 1997). The picture, however, is blurred because some studies reported a statistically significant motor illusion (Aglioti et al., 1995; Daprati & Gentilucci, 1997), whereas in other studies the effect of the illusion on grasping failed to reach signifi-

Address correspondence to Volker Franz, Max-Planck-Institut für Biologische Kybernetik, Spemannstr. 38, D-72076 Tübingen, Germany; e-mail: volker.franz@tuebingen.mpg.de.

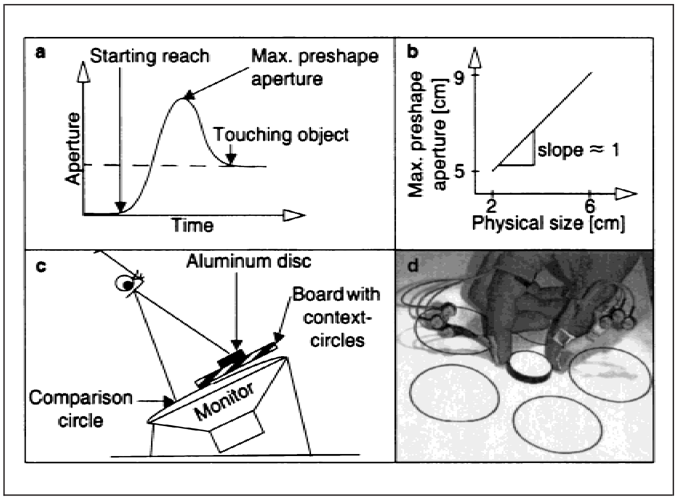

Fig. 1. Maximum preshape aperture and the apparatus of Experiment 1. The typical time course of the aperture between index finger and thumb during the transport component of a prehension movement is shown in (a). The maximum preshape aperture is linearly related to the physical size of the object (b). In Experiment 1, subjects viewed a board with either large or small context circles drawn on it (c). An aluminum disc was positioned in the center of the context circles. In the grasping task, subjects grasped the disc; in the perceptual task, they adjusted a comparison circle displayed on the monitor to match the size of the aluminum disc. The illustration in (d) shows a subject grasping the aluminum disc.

cance (Brenner & Smeets, 1996; Haffenden & Goodale, 1998). Using a larger sample size, we tried to settle this issue.

Second, the original studies compared the influence of the illusion on the perceptual measures directly with the influence on maximum preshape aperture. For example, Aglioti et al. (1995) found an influence of 2.5 mm on perception and of 1.6 mm on grasping. Because of the statistically significant difference between these values ($p < .02$, $N = 14$), they concluded that the influence of the illusion on the motor system was dissociated from the influence on perception. However, given that the function relating perceived size to physical size and the function relating maximum preshape aperture to physical size are both linear, this conclusion is valid only if these functions also have the same slopes. For example, if maximum preshape aperture depends on physical size with a slope of 1/2 and perceived size has a slope of 1, even the common-representation model would predict a motor illusion of only half the size of the perceptual illusion. To obtain a good estimate of these slopes, we used a wider range of disc sizes.

Third, in order to make the perceptual and motor tasks as similar as possible, we presented only one Ebbinghaus figure at a time: A central disc was surrounded by either large or small context circles (single-context versions, Fig. 2a). In the perceptual task, subjects indicated the size of the central disc by adjusting the radius of an isolated circle displayed on a monitor (Coren & Girgus, 1972; Pressey, 1977). In the grasping task, subjects grasped the central disc. In contrast, the original studies used the composite version of the illusion (Fig. 2b): Two Ebbinghaus figures with different context circles were presented simultaneously. In the perceptual task, subjects directly compared the two central discs. In the grasping task, however, subjects grasped only one of the discs on each trial. The overall effect of the illusion on grasping was then determined by adding the effects each context had on grasping. Note the asymmetry in this procedure. In grasping, subjects operated on only one Ebbinghaus figure at a time; in the perceptual task, they operated on both figures simultaneously. A perceptual task that is more similar to the grasping task of the original studies is shown in Figure 2c: On each trial, subjects compare an isolated circle with one of the central discs, and the overall effect of the illusion on perception is then determined by adding the effects of the two separate comparisons. In using the direct comparison instead, the original studies implicitly relied on an additivity assumption: It was assumed that the perceptual effects of the two Ebbinghaus figures simply add up to yield the effect obtained by the direct comparison. In our Experiments 2 and 3, we tested this assumption and found that it is not correct. A direct comparison between two Ebbinghaus figures (Fig. 2b) yields a larger effect than if the perceptual effects are determined for each figure separately and then added (Fig. 2c). Using this simplified and improved

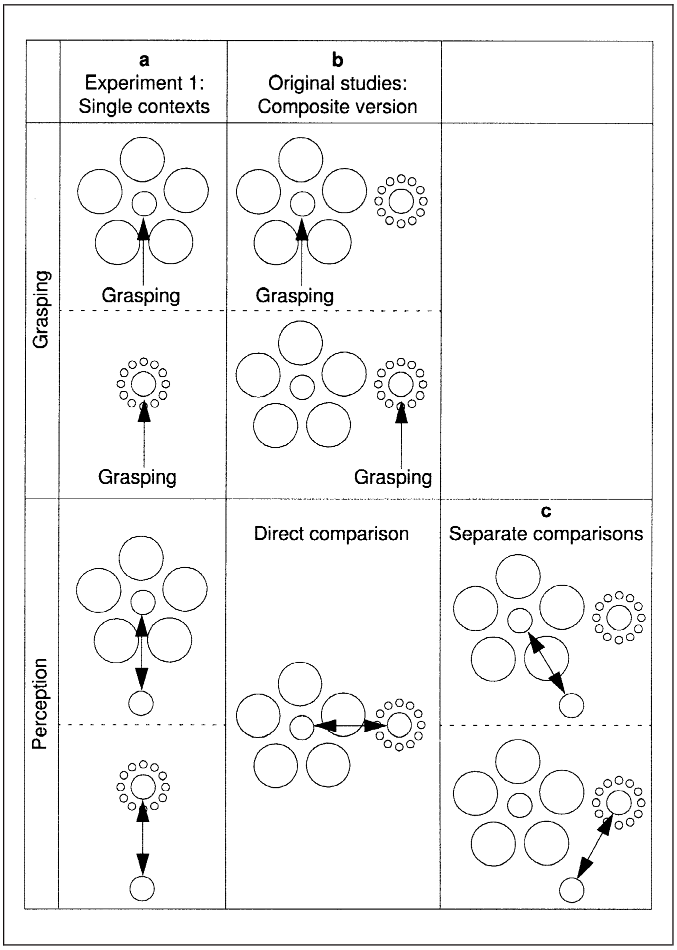

Fig. 2. Stimuli used in our experiments and in the original studies. In Experiment 1 (a), subjects operated on only one Ebbinghaus figure at a time (single-context versions). In the original studies (b), asymmetric measures were used: To perform the grasping task, subjects had to calculate only the size of one of the central discs. In the perceptual task, subjects had to compare the two central discs directly, both being subjected to the illusion at the same time. A perceptual task that is more similar to the grasping task of the original studies is shown in (c).

design, we investigated whether grasping is influenced by the Ebbinghaus illusion.

In all our experiments, subjects were students of the University of Tübingen, Germany. In return for their participation, they received a payment of 13 DM per hour. Subjects had normal or corrected-to-normal vision. Stimuli were chosen to be similar to the ones used in the original studies. In the large context, there were 5 circles 58 mm in diameter, and the centers of the circles were 118 mm apart. In the small context, there were 12 circles 10 mm in diameter, and the centers of the circles were 60 mm apart.

EXPERIMENT 1: GRASPING THE ILLUSION

Method

Twenty-six right-handed (Oldfield, 1971) subjects participated in Experiment 1. The isolated circle had a distance of 155 mm from the central disc. The central discs were 28, 31, 34, or 37 mm in diameter and 5 mm in height. The apparatus of Experiment 1 is shown in Figure 1c. Attempting to generate large effects of the illusion, we maximized the figural similarity (Coren & Miller, 1974) between the three-dimensional central disc and the two-dimensional context circles. This was achieved by minimizing shadows and having subjects view the disc from above. Subjects sat on a stool at a viewing distance of approximately 65 cm.

Subjects wore liquid-crystal shutter glasses that were opaque while the stimuli for each trial were set up. After this, the glasses became transparent. In the grasping task, subjects grasped the central disc with their right hand. As soon as they started to move their hand, the glasses became opaque again. Therefore, the subjects could see neither their hand nor the stimulus during grasping (open-loop condition; Haffenden & Goodale, 1998; Jeannerod, 1981; Post & Welch, 1996). The mean presentation time of the stimuli was 825 ms. The grasp trajectory was recorded using an Optotrak™ system: Three infrared light-emitting diodes were attached to the thumb and index finger (Fig. 1d), and the maximum preshape aperture between the finger tips was calculated for each grasp. In the perceptual task, subjects adjusted the comparison circle displayed on the monitor to match the size of the central disc. After they finished their adjustment, the glasses became opaque again. Each subject performed 72 grasps and 24 adjustments.

Results and Discussion

Results of Experiment 1 are shown in Figure 3. Analyses of variance revealed highly significant effects of the illusion on perception, $F(1, 25) = 144$, $p < .001$, and on grasping, $F(1, 25) = 15.2$, $p = .001$. Regression analyses showed that maximum preshape aperture and perceived size were linearly related to physical size. The slopes for perception ($s = 1.10 \pm 0.01$) and for grasping ($s = 1.12 \pm 0.06$) were similar, $t(25) = 0.35$, $p = .73$. As the reasoning in the introduction indicates, this finding allows a comparison of the illusion's effects. To this end, we pooled effects across all disc sizes (Fig. 4a). The magnitudes of the pooled effects in perception and grasping were equal, $t(25) = 0.07$, $p = .94$, and were in a range typically found for the Ebbinghaus illusion (Coren & Girgus, 1972; Coren & Miller, 1974). These results clearly contradict the notion that the effects of the Ebbinghaus illusion are dissociated between action and perception.

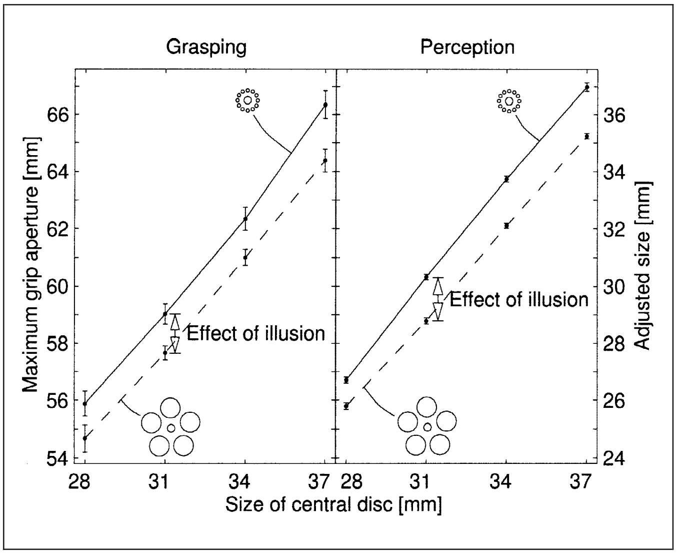

Fig. 3. Effects of the Ebbinghaus illusion on size perception and on maximum preshape aperture for each central-disc size in Experiment 1. Solid lines represent data for small context circles; dashed lines represent data for large context circles. Effects of the illusion are indicated by arrows. (Overall illusion effects pooled over all disc sizes are shown in Fig. 4a.) Error bars depict ± 1 standard error of the mean (data are normalized to account for absolute differences in aperture sizes between subjects).

An additional test for the common-representation model is the prediction that subjects showing a large perceptual illusion should also show a large motor illusion. Our data confirm this: We found a significant correlation of $\rho = .34$, $t(24) = 1.76$, $p < .05$. To evaluate the size of this correlation, we calculated and simulated a strong model of perception-action coupling: We assumed that for each subject the motor illusion equals the perceptual illusion—except that there is added noise in the motor system (this assumption is necessary to account for the larger variance of the motor data). Given this model, the expected correlation between the perceptual illusion and the motor illusion equals the ratio of their standard deviations. For our data, this results in $\rho = .32$ ($\sigma_P/\sigma_M = 0.62/1.93$). That is, we found exactly the correlation that is predicted by a strong model of perception-action coupling. The model also predicts that the motor illusion is related to the perceptual illusion with a slope of 1 and an intercept of 0. Again, the data agree well with this prediction: The slope was 1.06 and the intercept was −0.07.

But why, then, did the original studies find a difference between action and perception? A comparison shows that the grasping effects were similar in our Experiment 1 and in the original studies (Fig. 4). Only the perceptual effects were larger in the original studies. We hypothesized that this enhancement was due to the direct comparison in the perceptual tasks of the original studies (requiring the additivity assumption). To test this hypothesis, we conducted two perceptual experiments.

EXPERIMENT 2: TEST FOR ADDITIVITY

In Experiment 2, we measured the perceptual effects of the two single-context versions (as in Experiment 1; Fig. 2a), of a direct

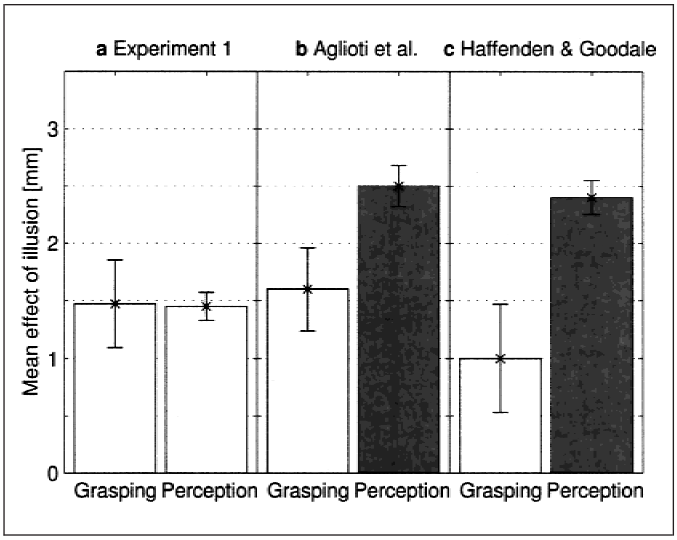

Fig. 4. Overall effects of the Ebbinghaus illusion on size perception and on maximum preshape aperture in Experiment 1 and in the original studies. Shaded bars indicate conditions in which a direct comparison between two Ebbinghaus figures was required. In Experiment 1, the illusion affected grasping just as much as perception (a). In the study by Aglioti, DeSouza, and Goodale (1995), the illusion affected grasping significantly less than perception (b). Haffenden and Goodale (1998) replicated the findings of Aglioti et al. (c). Error bars depict ± 1 standard error of the mean.

comparison (as in the original studies; Fig. 2b), and of two separate comparisons (our suggestion for a better perceptual measure in the original studies; Fig. 2c). If the additivity assumption of the original studies holds, then the direct comparison should yield an effect similar to the sum of the effects of the two separate comparisons.

Method

Eighteen subjects participated in Experiment 2. The central circles had diameters of 28, 31, and 34 mm. In the composite version, the centers of the central circles were 140 mm apart. The isolated circle had a distance of 140 mm from the central circle. Stimuli were identical to those used in Experiment 1 in all other respects, except that they were presented on a computer monitor. The central element was therefore a two-dimensional circle and no longer a three-dimensional disc. This increased figural similarity between central element and context elements, and therefore increased the magnitude of the illusion slightly (Coren & Miller, 1974).

All three possibilities for assessing the perceptual effect of the illusion were employed, as shown in the lower part of Figure 2. In the single-context condition, subjects adjusted the size of an isolated circle to match the size of the central circle in one Ebbinghaus figure. The effects of the large context circles and of the small context circles were added to obtain an estimate of the illusion strength. In the direct-comparison condition, subjects adjusted the central circles of the two Ebbinghaus figures simultaneously. The difference between the two central circles that was needed for them to be perceived as equal in size was used as the measure of illusion strength. In the separate-comparison condition, subjects viewed both Ebbinghaus figures, but adjusted the isolated circle to match the size of only one of the central circles. The illusion strength was calculated the same way as in the single-context condition. Each subject performed a total of 75 adjustments.

Results and Discussion

Figure 5a shows the illusion's effects in Experiment 2. Single-context versions and separate comparisons showed similar effects, $t(17) = 0.99$, $p = .34$; the effect of the direct comparison was larger than the sum of the effects in the two separate comparisons, $t(17) = 2.27$, $p = .04$, and also larger than the sum of the effects in the single-context versions, $t(17) = 3.68$, $p = .002$. That is, the direct comparison yielded a larger perceptual effect than the sum of the effects in the separate comparisons. This failure of additivity contradicts the original studies' implicit assumption that the perceptual effects of the two Ebbinghaus figures simply add up to yield the effect obtained by the direct comparison. Results also show that the effect in the separate comparisons is similar to the effect in the single-context versions. Given that the motor illusions in all studies were similar to the perceptual illusion in the single-context versions, this means that the additivity failure can account for the differences found between perception and grasping in the original studies.

These results show an interesting nonadditive effect in the Ebbinghaus illusion: If subjects directly compare two Ebbinghaus figures, they experience a larger-size illusion than is predicted by the sum of the size illusions experienced in each figure separately. Interestingly, most quantitative research on the Ebbinghaus illusion has been based on the single-context versions, whereas qualitative demonstrations of the illusion usually employ a direct comparison in the composite version and therefore exhibit an effect that is about 50% larger (Coren & Girgus, 1972, 1978).

The additivity failure indicates that the perceptual task and the

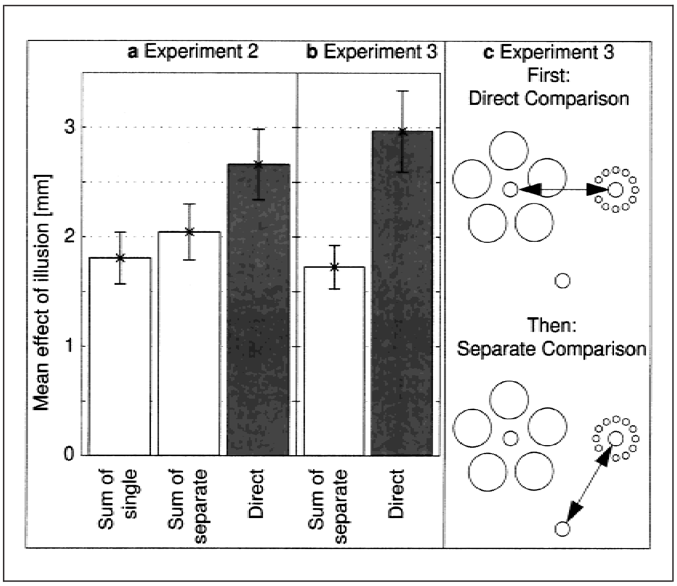

Fig. 5. Effects of the Ebbinghaus illusion in Experiments 2 (a) and 3 (b) and illustration of a trial in Experiment 3 (c). Error bars depict ±1 standard error of the mean.

grasping task were not appropriately matched in the original studies (Aglioti et al., 1995; Haffenden & Goodale, 1998). Therefore, one cannot conclude that the differences between perception and grasping in those studies were due to a dissociation between perception and action.

One possible objection to our argument is that in the original studies subjects had to directly compare the two central discs immediately before grasping (this was done as a control for the perceptual effect). Could this direct comparison, in which subjects were forced to attend to both central discs, induce additivity? We tested this possibility in Experiment 3.

EXPERIMENT 3: DID ATTENTION INDUCE ADDITIVITY?

In Experiment 3, we approximated the succession of the perceptual task and of the grasping task in the original studies more closely. Subjects first compared the two central circles directly and then, immediately afterward, compared one of the central circles to the isolated circle (Fig. 5c). Both comparisons were performed within 1 s, which is similar to the mean onset time for grasping in the study by Aglioti et al. (1995). If the first comparison induces additivity, then the effect in the direct comparison would be expected to equal the sum of the effects in the separate comparisons.

Method

Twelve subjects participated in Experiment 3. The central circles had diameters of 31, 32, 33, and 34 mm. In all other respects, the stimuli and the apparatus were identical to those of Experiment 2. A typical trial of Experiment 3 is shown in Figure 5c. Subjects performed the direct comparison and the separate comparison in direct succession, within 1 s. Because this short time interval did not allow an adjustment procedure, constant stimuli and two two-alternative forced-choice tasks were used. Each subject compared 462 configurations to complete the psychometric functions.

Results and Discussion

Results are shown in Figure 5b. As in Experiment 2, the effect of the illusion in the direct comparison was significantly larger than the sum of the effects in the two separate comparisons, $t(11) = 3.45$, $p = .006$. This result indicates that additivity cannot be induced by an immediate succession of direct comparison and separate comparison.

GENERAL DISCUSSION

If perceptual and motor tasks are carefully matched, there are strikingly similar effects of the Ebbinghaus illusion on perceived size and on maximum preshape aperture. In Experiment 1, we replicated the influence of the Ebbinghaus illusion on grasping found in the original studies of Aglioti et al. (1995) and Haffenden and Goodale (1998). However, our Experiments 2 and 3 show that the larger perceptual effect in the original studies is likely due to an additivity failure that selectively enhanced this effect.

To discuss our results in a more general context, we should mention that Haffenden and Goodale (1998) not only replicated the study of Aglioti et al. (1995), but also employed an additional task. Subjects estimated the size of one of the central discs using their thumb and index finger (without seeing their hand). The authors interpreted this manual estimation task as a perceptual measure. They found a significantly larger influence of the illusion for manual estimation (4.2 ± 0.97 mm) than for grasping (Fig. 4c). However, the effect on manual estimation was also larger than the effect on the classic perceptual measure (Fig. 4c). This difference is even more pronounced if one corrects the classic perceptual measure for the nonadditivity. Given that classic perceptual measures are much better understood (e.g., Coren & Girgus, 1972), it does not seem appropriate to infer a dissociation between perception and action based on the manual estimation task alone.

Our findings also have implications for studies investigating other visual illusions. Brenner and Smeets (1996) used the Aglioti paradigm to investigate the Ponzo illusion and found a smaller influence of the illusion on grasping (0.3 mm) than on perception (0.8 mm). This difference again might be due to a failure of additivity, as we found with the Ebbinghaus illusion.

In conclusion, the Ebbinghaus illusion does not provide evidence for different processing mechanisms for perception and action. To the contrary, our results strongly suggest that in the illusion the same internal representation is used for perception and for grasping. This outcome contradicts the predictions of the perception-versus-action hypothesis of Milner and Goodale (1995) and removes one critical piece of evidence that is usually counted in favor of this theory (e.g., Jackson & Husain, 1997).

Acknowledgments—We wish to thank M.S. Banks and M. Jeannerod for helpful comments on earlier versions of the manuscript. Also, we wish to thank M.A. Goodale and A.M. Haffenden for long and fruitful discussions about our opposite views. This work was supported by a grant from the Deutsche Forschungsgemeinschaft and by the Max-Planck Society.

REFERENCES

Aglioti, S., DeSouza, J.F.X., & Goodale, M.A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, *5*, 679–685.

Brenner, E., & Smeets, J.B.J. (1996). Size illusion influences how we lift but not how we grasp an object. *Experimental Brain Research*, *111*, 473–476.

Bridgeman, B., Lewis, S., Heit, G., & Nagle, M. (1979). Relation between cognitive and motor-oriented systems of visual position perception. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 692–700.

Coren, S., & Girgus, J.S. (1972). A comparison of five methods of illusion measurement. *Behavior Research Methods & Instrumentation*, *4*, 240–244.

Coren, S., & Girgus, J.S. (1978). *Seeing is deceiving: The psychology of visual illusions*. Hillsdale, NJ: Erlbaum.

Coren, S., & Miller, J. (1974). Size contrast as a function of figural similarity. *Perception & Psychophysics*, *16*, 355–357.

Daprati, E., & Gentilucci, M. (1997). Grasping an illusion. *Neuropsychologia*, *35*, 1577–1582.

Goodale, M.A., Meenan, J.P., Bülthoff, H.H., Nicolle, D.A., Murphy, K.J., & Carolynn, I.R. (1994). Separate neural pathways for the visual analysis of object shape in perception and prehension. *Current Biology*, *4*, 604–610.

Goodale, M.A., & Milner, A.D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, *15*, 97–112.

Goodale, M.A., Milner, A.D., Jakobson, L.S., & Carey, D.P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, *349*, 154–156.

Goodale, M.A., Pélisson, D., & Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, *320*, 748–750.

Haffenden, A.M., & Goodale, M.A. (1998). The effect of pictorial illusion on prehension and perception. *Journal of Cognitive Neuroscience*, *10*, 122–136.

Hansen, R.M., & Skavenski, A.A. (1985). Accuracy of spatial localization near the time of a saccadic eye movement. *Vision Research*, *25*, 1077–1082.

Jackson, S.R., & Husain, M. (1997). Visual control of hand action. *Trends in Cognitive Sciences*, *8*(1), 310–317.

Jeannerod, M. (1981). Intersegmental coordination during reaching at natural visual objects. In J. Long & A. Baddeley (Eds.), *Attention and performance* (Vol. 9, pp. 153–168). Hillsdale, NJ: Erlbaum.

Jeannerod, M. (1984). The timing of natural prehension movements. *Journal of Motor Behavior*, *16*, 235–254.

Milner, A.D., & Goodale, M.A. (1995). *The visual brain in action*. Oxford, England: Oxford University Press.

Oldfield, R.C. (1971). The assessment and analysis of handedness: The Edinburgh Inventory. *Neuropsychologia*, *9*, 97–113.

Post, R.B., & Welch, R.B. (1996). Is there dissociation of perceptual and motor responses to figural illusions? *Perception*, *25*, 569–581.

Pressey, A.W. (1977). Measuring the Titchener circles and Delboeuf illusions with the method of adjustment. *Bulletin of the Psychonomic Society*, *10*, 118–120.

Ungerleider, L.G., & Mishkin, M. (1982). Two cortical visual systems. In D.J. Ingle, M.A. Goodale, & R.J.W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549–586). Cambridge, MA: MIT Press.

(Received 11/30/98; Revision accepted 3/30/99)