

$$\hat{S} = \sum_i w_i \hat{S}_i \quad \text{with} \quad w_i = \frac{1/\sigma_i^2}{\sum_j 1/\sigma_j^2} \quad (2)$$

Thus, the MLE rule states that the optimal means of estimation (in the sense of producing the lowest-variance estimate) is to add the sensor estimates weighted by their normalized reciprocal variances^{8–15}. If the MLE rule is used to combine visual and haptic estimates, \hat{S}_V and \hat{S}_H , the variance of the final (visual–haptic) estimate, \hat{S} , is

$$\sigma_{VH}^2 = \frac{\sigma_V^2 \sigma_H^2}{\sigma_V^2 + \sigma_H^2} \quad (3)$$

Thus, the final estimate has lower variance than either the visual or the haptic estimator. Implementation of MLE integration is shown for two hypothetical cases in Fig. 1.

We examined visual–haptic integration quantitatively to determine whether human performance is optimal. Observers looked at and/or felt a raised ridge (Fig. 2) and judged its height (vertical extent). To work out the predictions of the MLE rule, we first determined the variances of the visual and haptic height estimates (within-modality) by conducting discrimination experiments. In the haptic-alone experiment, observers indicated which of two sequentially presented ridges was taller from haptic information alone; in the visual-alone experiment, they did the same from visual information alone. There were four conditions in the visual experiment that differed in the amount of noise in the stimulus (see Methods). By adding noise we made the visually specified height less reliable.

Visual-alone and haptic-alone discrimination data are shown in Fig. 3a. The proportion of trials in which the observer indicated that the comparison stimulus (variable height) appeared taller than the standard stimulus (fixed height of 55 mm) is plotted as a function of the height of the comparison stimulus. The dashed red line and

symbols represent the haptic discrimination data, and the solid blue curves with open symbols represent the visual data for the four levels of noise. These psychometric functions were well fit by cumulative gaussian functions.

The discrimination threshold is defined as the difference between the point of subjective equality (PSE) and the height of the comparison stimulus when it is judged taller than the standard stimulus 84% of the time. The 84% point corresponds to $\sqrt{2}$ times the standard deviation of the underlying estimator. The haptic discrimination threshold was roughly 0.085 times the average ridge height (which was 55 mm). As the noise increased from 0 to 200%, the visual discrimination thresholds increased from 0.04 to 0.2 times the average height. Thus, when the visual noise was 0%, the visual discrimination threshold was roughly half the haptic threshold; when the visual noise was 200%, the visual threshold was more than double the haptic threshold.

In the visual–haptic experiment, observers simultaneously looked at and felt two raised ridges that were presented sequentially. In one presentation the visually and haptically specified heights were equal (comparison stimulus); in the other presentation they differed (standard stimulus). The difference (Δ) in the specified heights was ± 6 , ± 3 or 0 mm (the average of S_H and S_V was 55 mm). For each Δ in the standard stimulus (randomly presented), the height of the comparison stimulus was varied randomly from trial to trial (47–63 mm). On each trial, the observer indicated which stimulus seemed taller. Figure 3b shows the proportion of trials in which the comparison stimulus was chosen as taller as a function of the height of the comparison stimulus. From these psychometric functions, we estimated the PSE—the comparison height appearing equal to the standard height—and the just-discriminable change in height (threshold).

Using the within-modality data, we can predict what an observer using MLE will do when presented visual and haptic information simultaneously, and we compare these predictions to the perfor-

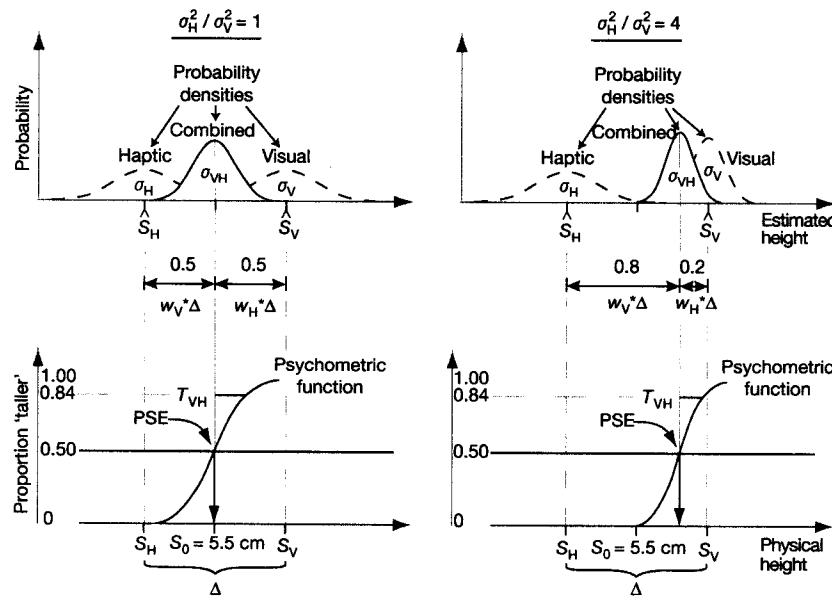


Figure 1 Maximum-likelihood estimation integration: two hypothetical situations. Visually and haptically specified heights differ by Δ . Dashed Gaussians in the top panels represent probability densities of the (unbiased) estimated height from visual and haptic assessment, and solid Gaussians represent probability densities for the combined estimate. On the left, the visual and haptic variances are equal ($\sigma_H^2/\sigma_V^2 = 1$) and both their weights are 0.5 (equation (2)). The mean of the combined probability density is therefore equal to the mean of the visual and haptic densities and the variance is reduced by half (equation (3)). If the observer bases judgements of relative height on the combined

probability density, the psychometric function would be a cumulative gaussian (bottom left) with a point of subjective equality (PSE) equal to the average of the visual and haptic heights of the standard stimulus. On the right, the haptic variance is four times the visual variance: $\sigma_H^2/\sigma_V^2 = 4$. By equation (2), the visual weight (w_V) is 0.8 and the haptic weight (w_H) is 0.2. Thus, the combined probability density is shifted towards the visual estimate and its variance is 0.8 times the visual variance (equation (3)). Accordingly, the psychometric function should be shifted so that the PSE is closer to the visual height of the standard stimulus.

mance in the visual–haptic experiment. First, we describe the analysis of the PSE data and predictions for the weights. From equation (2) and the relationship between threshold and estimator variance:

$$\frac{w_V}{w_H} = \frac{\sigma_H^2}{\sigma_V^2} = \frac{T_H^2}{T_V^2} \quad (4)$$

where T_H and T_V are the haptic and visual thresholds (84% points in Fig. 3a). Incorporating the normalization assumption ($w_V + w_H = 1$), the predicted weights for optimal integration are

$$w_V = \frac{T_H^2}{T_V^2 + T_H^2} \quad \text{and} \quad w_H = \frac{T_V^2}{T_V^2 + T_H^2} \quad (5)$$

The predicted visual weights are represented by the curve and shaded surround in Fig. 3c. The predicted weights vary significantly with the amount of visual noise in the stimulus: the visual weights are higher when the noise level is low, and lower when the noise level is high. Assuming that the visual and haptic estimators are on average unbiased ($\hat{S}_V = S_V$ and $\hat{S}_H = S_H$), the weights can be derived experimentally:

$$w_V = (\text{PSE} - S_H)/(S_V - S_H) \quad (6)$$

where PSE is the height of the comparison stimulus that matched the apparent height of the standard stimulus. The visually and haptically specified heights in the standard stimulus, S_V and S_H , are indicated on the right ordinate. Figure 3c shows that as the noise level was increased the visual weight decreased, and the PSE shifted from S_V towards S_H . Because the noise level varied randomly from trial to trial, the weights must have been set within the 1-s stimulus presentation. Below, we suggest a mechanism for such dynamic weight adjustment. In summary, the predicted and observed PSEs are similar, suggesting that humans do combine visual and haptic information in a fashion similar to MLE integration.

According to the MLE rule, the combined estimates should have lower variance, and therefore lower discrimination thresholds, than either the visual or haptic estimate alone (equation (3)). To derive predictions for the visual–haptic discrimination thresholds, we rewrite equation (3):

$$T_{VH}^2 = \frac{T_V^2 T_H^2}{T_V^2 + T_H^2} \Leftrightarrow \frac{1}{T_{VH}^2} = \frac{1}{T_V^2} + \frac{1}{T_H^2} \quad (7)$$

The predicted and observed thresholds are shown in Fig. 3d. The open symbols represent the visual-alone thresholds and the dashed line represents the haptic-alone threshold. The shaded area represents the predicted visual–haptic thresholds; they are always lower than the visual-alone and haptic-alone thresholds at the corresponding noise level. The filled purple symbols represent the observed visual–haptic discrimination thresholds; as noise level increases, the just-noticeable difference in height becomes greater. Most notably, the predicted and observed visual–haptic discrimination thresholds are similar. As with the PSE data, this indicates that human observers may combine visual and haptic information in a manner similar to MLE integration.

In summary, we found that height judgements were remarkably similar to those predicted by the MLE integrator. Thus, the nervous system seems to combine visual and haptic information in a fashion similar to the MLE rule: visual and haptic estimates are weighted according to their reciprocal variances (equation (2)). Naturally, it is important to determine whether this rule characterizes the estimation of other stimulus properties such as shape, depth, localization, roughness or compliance.

The relative contributions of vision and haptics to perceiving such object properties have been studied previously. For example, subjects have grasped a square while looking at it through a distorting lens that made it appear rectangular¹. The shape of the unified percept was determined almost completely by vision, so the phenomenon was called ‘visual capture’. Numerous studies have




Figure 2 Apparatus and stimuli. **a**, Observers viewed the reflection of the visual stimulus, presented on a cathode ray tube (CRT) binocularly in a mirror. CrystalEyes (StereoGraphics) liquid-crystal shutter glasses were used to present binocular disparity. The surfaces of the stimuli were perpendicular to the line of sight. A head and chin rest limited head movements. The right hand was beneath the mirror and could not be seen. The haptic stimulus was presented with two PHANTOM force-feedback devices, one each for the index finger and thumb. **b**, Stereograms representative of visual stimuli, which should be viewed by cross-fusing. Top row, stereogram has no noise and the horizontal bar is raised above the background. Bottom row, stereogram of bar and background contains noise, with random displacements of dots parallel to the line of sight.

replicated visual capture in shape and size perception^{3,16,17}, depth perception¹⁸ and localization^{2,19–21}. But visual capture does not occur in the perception of surface roughness^{6,22}; instead, perceived roughness is affected nearly equally by haptics and vision. Does a dynamic cue-combination rule, such as the one described here, determine the degree to which vision or haptics dominates? The statistically optimal means of combining visual and haptic information—the MLE rule—predicts that visual capture should occur whenever the visual estimate of a property has much less variance than that of the haptic estimate. Haptic capture should be observed when the reverse occurs. We observed behaviour like visual capture when the visual stimulus was noise-free, and behaviour similar to haptic capture when the visual stimulus was quite noisy (Fig. 3c).

In visual–haptic tasks, the MLE integrator described here always uses information from both sensory systems, so the combined percept will always reflect both sources of information. Different behaviour may be observed when the discrepancy between two information sources is large. With large discrepancies between information sources, the nervous system may exhibit robust behaviour¹⁰ in which a discrepant source is discounted^{9,23}. In robust estimation, the weights associated with different information sources are determined by more than the variances of the sensory estimators; they are also determined by the discrepancy between their estimates. In our experiments, the differences between the visually and haptically specified stimuli were never greater than 11%, and the visual–haptic data always exhibited an influence of both sensory systems. If the differences had been larger, we might

have observed discounting of one sense. It would be interesting to know whether such vetoing occurs when observers become aware of the conflict between visual and haptic inputs.

If the nervous system implements MLE integration, the weights must be proportional to the reciprocal variances of the probability densities associated with the visual and haptic estimates of the environmental property in question (equation (2) and Fig. 1). Of course, the variances change from one object property to the next (such as size, shape or roughness) and from one situation to another (for example, visual variance increases as the lighting is degraded). Does the nervous system need to calculate or learn the variances associated with the visual and haptic estimators for each property and situation to implement MLE integration? Although explicit calculation or learning may occur²⁴, there are plausible schemes in which explicit calculation of variances or weights is unnecessary.

Consider, for example, a population of visual and haptic neurons, each sensitive to a range of heights. Each neuron has a preferred height but also responds to other heights (that is, each neuron has a tuning function). If the visual input specifies the height clearly (that is, high contrast, noise-free, and so on), then the visual neurons preferring that height respond vigorously and those preferring other heights respond less, and the distribution of response across the population of visual neurons has a well defined peak. Assume that the distribution across the population of haptic neurons has a less well defined peak. Multiplication of these two distributions (point-by-point multiplication where the two populations are in registration according to the stimulus property being estimated)²⁵ yields a

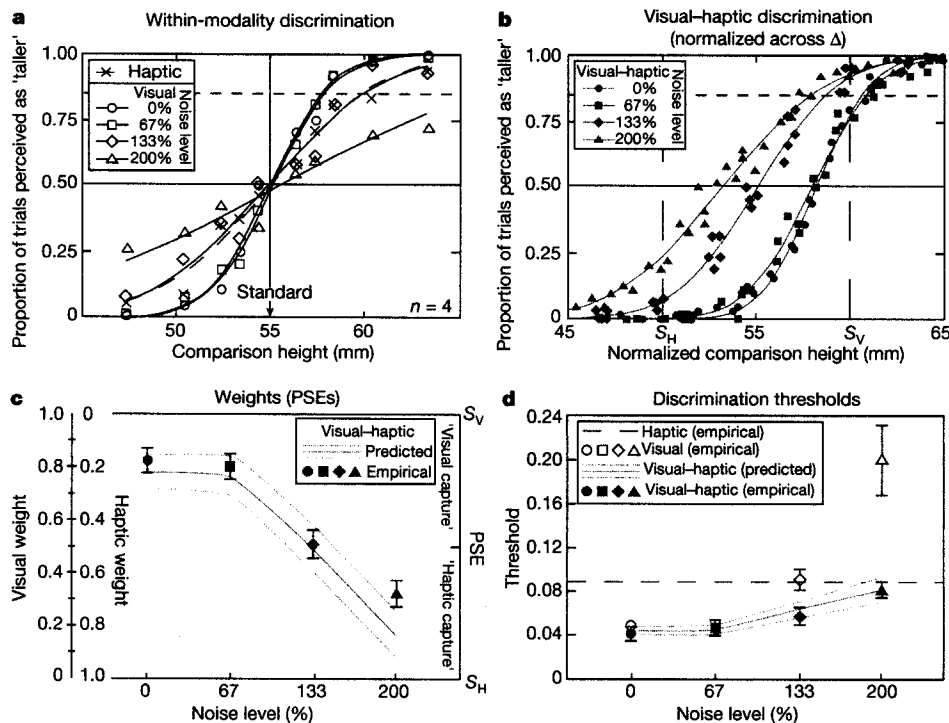


Figure 3 Predictions and experimental data. **a**, Within-modality discrimination. Proportion of trials in which a comparison was perceived as taller than the standard stimulus is plotted against the height of the comparison stimulus. Data were averaged for observers. Height of the standard stimulus was 55 mm. Haptic discrimination data are represented by red crosses and the dashed curve (best-fitting cumulative gaussian); visual discrimination data are represented by blue curves, which correspond to four noise levels. **b**, Visual–haptic discrimination. The average height of visual–haptic standard stimulus was 55 mm; the height difference, Δ , varied from -6 to $+6$ mm. To plot the data on the same coordinates, the psychometric functions for each Δ were shifted laterally by $w_V\Delta/2$ (w_V is obtained from the regression of PSE against Δ). Purple curves represent the different visual noise levels. **c**, Weights and PSEs. Abscissa represents the noise level, left

ordinate represents visual weight (w_V ; haptic weight is $1 - w_V$) and right ordinate represents the PSEs relative to S_V and S_H . Purple symbols represent observed visual weights obtained from regression analysis of PSEs (equation (6)) across Δ . Shaded area represents predicted weights expected from within-modality discrimination (**a**; equation (5)); its height represents predicted errors given the standard errors of the within-modality discrimination. **d**, Combined and within-modality discrimination thresholds. Just-noticeable differences in height are plotted against noise level. Thresholds are taken from psychometric functions in **a** and **b**. Dashed red line represents haptic-alone threshold; open blue symbols represent visual-alone thresholds; filled purple symbols represent combined visual–haptic thresholds. Shaded area represents predicted visual–haptic thresholds (equation (7)).

peak response closer to the visual than the haptic peak, as with MLE integration. If degrading the visual input causes the response distribution of the visual neurons to spread, then multiplication of the visual and haptic distributions yields a peak closer to the haptic peak, again as in MLE integration. Thus, the estimator variances (and therefore the weights) do not have to be calculated explicitly: the behaviour of an MLE integrator might be achieved through interactions among populations of visual and haptic neurons. □

Methods

Stimuli

The stimulus was a horizontal bar raised 30 mm above a plane; the bar and plane were perpendicular to the line of sight (Fig. 2a). The width of the bar was 150 mm; its height varied but the average was $S_0 = 55$ mm. Observers viewed the bar binocularly and/or grasped it with the index finger and thumb in order to estimate its height (Fig. 2a). Its vertical position was varied randomly from trial to trial.

The haptic stimulus was generated using two PHANTOM (SensAble Technologies) force-feedback devices (Fig. 2a), one each for the index finger and thumb. Finger and thumb movements had all six degrees of freedom in the 20-cm³ workspace. The three-dimensional positions of the tips of the finger and thumb were monitored, and appropriate forces were applied when the tips reached the positions of the simulated haptic objects. PHANTOMS compellingly simulate haptic properties such as the size, shape and stiffness of the bar. The apparatus was calibrated to align the visual and haptic stimuli spatially.

The visual stimulus was a random-dot stereogram simulating the background plane and bar (Fig. 2). The dots subtended 8 arcmin at the 50-cm viewing distance. Dot density was roughly 9 dots per degree². New dots were displayed with each presentation. The positions of the finger and thumb (tracked by the PHANTOMS) were indicated by small three-dimensional markers; the markers were visible until the bar was touched. Noise was added to the visual display to vary its reliability. The noise was a random displacement of the dot depths in the stereogram (direction parallel to line of sight). The displacements were drawn from a random uniform distribution whose range was 0, 67, 133 or 200% of the 3-cm depth step between the bar and plane. The displacement of the dots is shown in Fig. 2. No noise was added to the haptic display.

We wanted the presentation times for the visual and haptic stimuli to be identical. In the vision-alone discrimination experiment, the standard and comparison stimuli were displayed for 1 s each. In the haptic-alone discrimination experiment, the haptic stimulus began when the thumb and finger both contacted the bar and ended after 1 s. In the visual-haptic trials, the visually specified bar did not appear until the bar was touched by both fingers simultaneously and the visual and haptic stimuli were extinguished after 1 s.

Procedure

Within-modality discrimination was measured in a two-interval, forced-choice scheme. Each trial consisted of the sequential (visual or haptic) presentation of two bars. In the standard interval, the bar was always 55 mm tall and in the comparison interval, it was shorter or taller than 55 mm. The standard and comparison stimuli were randomly assigned to the first or second interval. The observer indicated the interval containing the apparently taller stimulus. The comparison height was varied according to the method of constant stimuli. We plotted psychometric functions, that is the proportion of trials in which the comparison was perceived as taller than the standard against the comparison height. From these functions, we could determine PSEs and discrimination thresholds. Half of the vision-alone and haptic-alone discrimination experiments were conducted (in random order) before the visual-haptic experiment and the other half were conducted after.

In the visual-haptic experiment, observers were again presented two stimuli sequentially: the standard stimulus (visually and haptically specified heights differed by Δ) and the comparison stimulus (visually and haptically specified heights were the same). The average height of the standard stimulus was always 55 mm and the visual-haptic difference (Δ) ranged from -6 to +6 mm. The height of the comparison varied. The observer indicated the interval containing the apparently taller stimulus. Trials with different conflicts were presented in random order to prevent visuomotor adaptation.

No feedback was given in any of the experiments. Four observers (aged 22 to 33), naive to the experimental purpose, participated. They were right handed and had normal or corrected-to-normal vision. The observers were questioned at the end: none of them had noticed the conflicts between the visually and haptically specified heights.

Received 25 July; accepted 15 November 2001.

1. Rock, I. & Victor, J. Vision and touch: An experimentally created conflict between the two senses. *Science* **143**, 594–596 (1964).
2. Hay, J. C., Pick, H. L. Jr & Ikeda, K. Visual capture produced by prism spectacles. *Psychonomic Sci.* **2**, 215–216 (1965).
3. Warren, D. H. & Rossano, M. J. in *The Psychology of Touch* (eds Heller, M. A. & Schiff, W.) 119–137 (Erlbaum, Hillsdale, New Jersey, 1991).
4. Power, R. P. The dominance of touch by vision: Sometimes incomplete. *Perception* **9**, 457–466 (1980).
5. Welch, R. B. & Warren, D. H. Immediate perceptual response to intersensory discrepancy. *Psychol. Bull.* **88**, 638–667 (1980).

6. Lederman, S. J. & Abbott, S. G. Texture perception: Studies of intersensory organization using a discrepancy paradigm, and visual versus tactual psychophysics. *J. Exp. Psychol. Hum. Percept. Perform.* **7**, 902–915 (1981).
7. Heller, M. A. Haptic dominance in form perception with blurred vision. *Perception* **12**, 607–613 (1983).
8. Clark, J. J. & Yuille, A. L. *Data Fusion for Sensory Information Processing Systems* (Kluwer Academic, Boston, 1990).
9. Blake, A., Bülthoff, H. H. & Sheinberg, D. Shape from texture: Ideal observer and human psychophysics. *Vision Res.* **33**, 1723–1737 (1993).
10. Landy, M. S., Maloney, L. T., Johnston, E. B. & Young, M. Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Res.* **35**, 389–412 (1995).
11. Gharamani, Z., Wolpert, D. M. & Jordan, M. I. in *Self-organization, Computational Maps, and Motor Control* (eds Morasso, P. G. & Sanguineti, V.) 117–147 (Elsevier, Amsterdam, 1997).
12. Knill, D. C. Discrimination of planar surface slant from texture: Human and ideal observers compared. *Vision Res.* **38**, 1683–1697 (1998).
13. Backus, B. T. & Banks, M. S. Estimator reliability and distance scaling in stereoscopic slant perception. *Perception* **28**, 417–442 (1999).
14. van Beers, R. J., Sittig, A. C. & Denier van der Gon, J. J. Integration of proprioceptive and visual position information: An experimentally supported model. *J. Neurophysiol.* **81**, 1355–1364 (1999).
15. Schrater, P. R. & Kersten, D. How optimal depth cue integration depends on the task. *Int. J. Comp. Vis.* **40**, 71–89 (2000).
16. Gibson, J. J. Adaptation, after-effect, and contrast in the perception of curved lines. *J. Exp. Psychol.* **16**, 1–31 (1933).
17. Festinger, L., Burnham, C. A., Ono, H. & Bamber, D. Efference and the conscious experience of perception. *J. Exp. Psychol.* **74** (4), 1–36 (1967).
18. Singer, G. & Day, R. H. Visual capture of haptically judged depth. *Percept. Psychophys.* **5**, 315–316 (1969).
19. Tastevin, J. En partant de l'expérience d'Aristote. *L'Encephale* **1**, 57–84 (1937).
20. Mon-Williams, M., Wann, J. P., Jenkinson, M. & Rushton, K. Synaesthesia in the normal limb. *Proc. R. Soc. Lond. B* **264**, 1007–1010 (1997).
21. Pavani, F., Spence, C. & Driver, J. Visual capture of touch: out-of-the-body experiences with rubber gloves. *Psycholog. Sci.* **11**, 353–359 (2000).
22. Heller, M. A. Visual and tactual texture perception: Intersensory cooperation. *Percept. Psychophys.* **31**, 339–344 (1982).
23. Banks, M. S. & Backus, B. T. Extra-retinal and perspective cues cause the small range of the induced effect. *Vision Res.* **38**, 187–194 (1998).
24. Ernst, M. O., Banks, M. S. & Bülthoff, H. H. Touch can change visual slant perception. *Nature Neurosci.* **3**, 69–73 (2000).
25. Peña, J. L. & Konishi, M. Auditory spatial receptive fields created by multiplication. *Science* **292**, 249–252 (2001).

Acknowledgements

We thank M. Landy for comments on the manuscript; and H. Ernst, X. Moncada, C. Alderson and S. Kashiwada for participating as observers. This research was supported by research grants from Air Force Office of Scientific Research and the National Institutes of Health, and by an equipment grant from Silicon Graphics.

Competing interests statement

The authors declare that they have no competing financial interests.

Correspondence and requests for materials should be addressed to M.O.E. (e-mail: marc.ernst@tuebingen.mpg.de).