

Prior knowledge of illumination for 3D perception in the human brain

Peggy Gerardin^{a,b,c}, Zoe Kourtzi^a, and Pascal Mamassian^{b,1}

^aSchool of Psychology, University of Birmingham, Birmingham B15 2TT, United Kingdom; ^bLaboratoire Psychologie de la Perception, Centre National de la Recherche Scientifique, Université Paris Descartes, 75006 Paris, France; and ^cLaboratoire Espace et Action, Institut National de la Santé et de la Recherche Médicale, 69676 Bron, France

Edited by Wilson S. Geisler, The University of Texas, Austin, TX, and approved August 6, 2010 (received for review May 10, 2010)

In perceiving 3D shape from ambiguous shading patterns, humans use the prior knowledge that the light is located above their head and slightly to the left. Although this observation has fascinated scientists and artists for a long time, the neural basis of this “light from above left” preference for the interpretation of 3D shape remains largely unexplored. Combining behavioral and functional MRI measurements coupled with multivoxel pattern analysis, we show that activations in early visual areas predict best the light source direction irrespective of the perceived shape, but activations in higher occipitotemporal and parietal areas predict better the perceived 3D shape irrespective of the light direction. These findings demonstrate that illumination is processed earlier than the representation of 3D shape in the visual system. In contrast to previous suggestions, we propose that prior knowledge about illumination is processed in a bottom-up manner and influences the interpretation of 3D structure at higher stages of processing.

functional MRI | multi-voxel pattern analysis | shape-from-shading | visual cortex | bottom-up processing

Although the perception of 3D shape is critically important for actions and interactions in the environments we inhabit, most depth cues are ambiguous. As a result, the brain requires additional information based on previous experience with the environment to infer 3D shape from depth cues. In particular, the inference of 3D shape from shading patterns (i.e., using image luminance intensity variations to derive the shape of a surface) relies on the assumption that the scene illuminant is above our heads and slightly to the left (1, 2). Understanding the illumination of a visual scene has fascinated artists and scientists for a long time (3, 4), but the neural basis of this “light from above left” preference for the interpretation of 3D shape remains largely unexplored. Although the light-from-above preference is consistent with an ecological explanation, the left bias remains entirely unexplained. More generally, the light-from-above preference provides a simple example of the way the brain represents prior knowledge and opens the door for the investigation of other types of prior knowledge related to our perception of the motion, shape, and color of objects.

Previous neurophysiological and imaging studies have implicated several brain regions in the processing of shape-from-shading: primary visual cortex (V1) (5–9), areas in the caudal inferior temporal gyrus (10) and the inferior parietal sulcus (11). However, the functions mediated by these different cortical areas may differ. In particular, interpreting shape-from-shading may involve at least two different stages of processing. At the first stage, the contrast polarity of edges in the image (dark or bright) is analyzed and related to the light direction so that left and right light directions can be discriminated. At the second stage, contrast edges are grouped together to form 3D shapes so that convex and concave shapes can be discriminated. For simplicity, we shall refer to these two stages as “Light” and “Shape” processing, respectively.

Here, we combine psychophysics and functional MRI (fMRI) measurements to dissociate the cortical areas engaged in the computations at the Light vs. the Shape stage of processing. We

use advanced fMRI analysis methods (multivoxel pattern analysis, MVPA) that allow us to evaluate whether small biases across voxels related to the preference of the underlying neural populations are statistically reliable (12–14). Using these methods, we test for fMRI sensitivity in discriminating Light and Shape across various visual, temporal, and parietal regions known to be engaged in the representation of shape from different depth cues (15–19). We then compare fMRI sensitivity for Light and Shape discrimination to that predicted from the behavioral data (20). Our fMRI results demonstrate that this left-light bias is processed at early stages of visual processing. In contrast, 3D shape perception is predicted from activations in parietal areas. In contrast to previous suggestions that prior knowledge always relates to top-down processes (e.g., 21), our findings support processing of some prior knowledge about the environment in a bottom-up manner.

Results

Psychophysics. Seven observers were presented with a series of images of shaded objects that belonged to eight image types (Fig. 1A, images “a” to “h”). These eight image types correspond to the interaction of two possible shapes (convex and concave rings) lit from one of four possible light directions (separated by 45°) (Fig. 1B). The ring was divided in eight equal sectors, all but one having the same shape (the odd sector is numbered 4 in image “a” of Fig. 1A). Observers were instructed to report the shape of the odd sector, convex or concave. The odd-shaped sector appeared in all six possible locations for each of the eight image types, so target location was not predictive of light or shape. Importantly, a concave ring lit from one direction produces the same image as a convex ring lit from the opposite direction. Fig. 1C shows the probability of the ring being perceived as convex (and the odd sector concave) for the eight image types. The ring was perceived as convex when the light was simulated from above (images “a,” “b,” “g,” and “h” in Fig. 1C) and concave otherwise. A repeated-measures ANOVA showed a significant main effect of light source position on the perceived shape [$F(7, 48) = 8.79, P < 0.0001$]. As expected from our previous behavioral work (20), the most ambiguous ring shape (the crossing points between convex and concave) did not occur for horizontal light directions: that is, halfway between images “b” and “c” and “f” and “g.” This observation illustrates a bias of the observers to prefer a light source located to the left of the vertical. To quantify this bias, we fitted the data in Fig. 1C with a scaled sinewave and used the phase of the sinewave as the estimate for the left bias (2). Averaging across all observers showed a bias to the left of the vertical equal to -22.3° .

Author contributions: P.G., Z.K., and P.M. designed research; P.G. performed research; Z.K. and P.M. contributed new reagents/analytic tools; P.G. and P.M. analyzed data; and P.G., Z.K., and P.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: pascal.mamassian@univ-paris5.fr.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1006285107/-DCSupplemental.

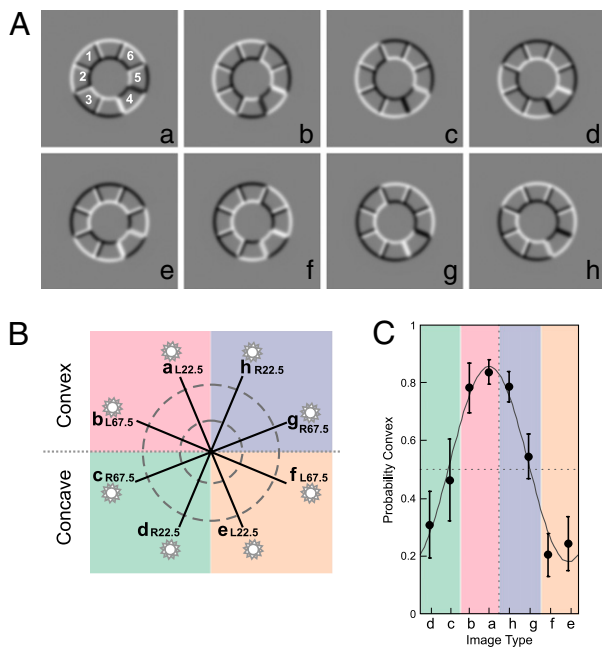


Fig. 1. Stimuli and behavioral performance. (A) Examples of stimuli. Each image is interpreted as a convex or concave ring lit from one of four light directions. All eight image types (a–h) are ambiguous. For instance, image “a” can be interpreted as a convex ring with a light source located above-left or as a concave ring with a light below-right. One of the elements of the ring, randomly chosen from one of six possible locations (numbered 1 through 6 in image “a”), has a shape opposite to that of the ring. In the behavioral task, observers were asked to report the perceived shape of this odd element. (B) Four classes of stimuli. To simplify the description of the stimuli in this article, we adopt the convention that the depicted shape of a stimulus is that consistent with a light coming from above. Following this convention, images “a,” “b,” “g,” and “h” will be referred to as convex rings and images “c,” “d,” “e,” and “f” as concave rings. The four main classes of stimuli are assigned different color codes: pink for convex shape lit from the left, green for concave-right, orange for concave-left, and blue for convex-right. (C) Behavioral performance in discriminating the shape of the odd element. The plot shows the probability that observers reported a convex ring (thus a concave odd element) as a function of light direction. In this plot, all six possible locations of the odd elements were pooled. The most ambiguous images were “c” and “g.” The solid line is the best fit of a scaled cosine function to illustrate the bias to above-left for the assumed light direction. Error bars are SEs across observers ($n = 7$).

Furthermore, to evaluate the statistical significance of the left bias, we compared the perceived shape of the ring when the light was simulated on the left versus when the light was on the right. That is, we compared significant differences in performance for pairs of images that were symmetric to the vertical meridian, namely “a,” “b,” “c,” and “d” vs. “h,” “g,” “f,” and “e” in Fig. 1.




Fig. 2. Shape-from-shading responsive regions. Group GLM map across subjects ($n = 7$) representing areas that were significantly more activated for shape-from-shading than scrambled stimuli [P (Bonferroni corrected) < 0.05]. The functional activations are superimposed on flattened cortical surfaces of the left and right hemispheres. The sulci are coded in darker gray than the gyri.

The first set of images were perceived significantly more often as convex (assuming light from above) compared with the second set [$t(6) = 7.2$, $P < 0.001$]. This difference is mostly because of the light directions that are close to the horizontal meridian (see also *SI Results* and Fig. S1).

fMRI Data: Shape-from-Shading Responsive Regions. For each individual participant we identified retinotopic, motion-related (V3B/KO, hMT+/V5) and shape-related (lateral occipital complex, LOC) areas based on standard procedures (see *SI Methods* for details and Figs. S2–S3). In addition to identifying regions involved in the processing of 3D shape, we compared activations [General Linear Modeling (GLM) analysis] for shape-from-shading ring stimuli to scrambled images of these stimuli. As shown in Fig. 2 (Table S1), we observed significantly stronger activations [P (Bonferroni corrected) < 0.05] for shape-from-shading stimuli in lateral occipital (LO), V3B/KO regions along the intraparietal sulcus (IPS) [ventral intraparietal sulcus (VIPS), parieto-occipital intraparietal sulcus (POIPS), dorsal intraparietal sulcus (DIPS)], the postcentral sulcus (left hemisphere) and ventral premotor region (right hemisphere). GLM analysis comparing responses to convex vs. concave shapes did not show any significant activations [P (Bonferroni corrected) < 0.05]. We then identified these shape-from-shading responsive regions in individual observers ($P < 0.05$ uncorrected) and used them as regions of interest (ROI) for the fMRI pattern classification analysis of light and shape processing (MVPA). This procedure ensured that the data used for the fMRI pattern classification were independent from the data used for the localization of ROI. Data from areas postcentral and ventral premotor were not further analyzed, as these areas were activated in less than three of the observers.

fMRI Multivoxel Pattern Classification: Discriminating Light and Shape Processing. Previous studies have shown that the perception of shape-from-shading involves a network of cortical areas (5–11). However, these studies have not decoupled the processing of Light from that of Shape. Here, Light processing refers to the ability to discriminate left from right light directions, whereas Shape processing refers to the ability to discriminate convex from concave shapes.

We used MVPA to test the extent to which neural populations in retinotopic areas, motion-related areas (V3B/KO, hMT+/V5), shape-related areas (LOC), and shape-from-shading responsive regions are involved in Light vs. Shape processing (see *Methods* for details). MVPA has been previously used successfully for decoding basic visual features (22–24), depth structure (25), and object categories (26–29) from fMRI data. We reasoned that activity from regions involved in Shape processing would contain information that distinguishes convex from concave shapes across different light directions. In contrast, we predicted that regions involved in Light processing would contain information that distinguishes left from right illuminations. It is important to note that we will evaluate the relative differences between classifiers across the four lighting direction conditions, rather than simply the ab-

solute performance of classifiers across ROIs. The advantage of this analysis is that it is independent of any variables that may contribute to the absolute classification values in each ROI (e.g., signal-to-noise ratio, partial volume effects). We conducted two main analyses for Shape and Light processing, respectively.

Shape classification. We conducted MVPA to test which cortical areas contain information about 3D shape (convex vs. concave). The analysis using the Shape classifiers resembles the task of the observers: we trained linear classifiers (SVM) to discriminate whether fMRI activity across voxel patterns relates to convex or concave rings. Four classifiers were used corresponding to four light directions (Fig. 3A). The first classifier was trained to distinguish convex from concave shapes when lit from 67.5° left: that is, to differentiate the shapes depicted in images “b” and “f” of Fig. 1A. The second classifier performed the same classification when shapes were lit from 22.5° left: that is, it differentiated images “a” and “e”. The remaining two classifiers performed the

same classification when shapes were lit from 22.5° right or 67.5° right (Fig. 3A).

What is the classification performance that we can expect for different light directions in cortical areas involved in 3D shape processing? Looking at the two images in a pair (e.g., “b” and “f”), we note that almost all white pixels from one image are black in the other. Therefore, there is some information in the image intensity that could support the discrimination of the stimuli; but importantly, this information is the same for all pairs of images (*SI Results*). In other words, a classifier that uses purely the image intensity to classify stimuli should have constant performance across pairs of stimuli (Fig. 3B).

In contrast, based on the behavioral data, we know that observers perceived image “a” as a convex ring much more frequently (about 60% more often) than image “e” (Fig. 1C). Therefore, we predict high classification accuracy for discriminating between these images from fMRI signals sensitive to the 3D shape (convex vs. concave). Conversely, we predict low classification accuracy for signals related to images “c” and “g” because observers perceived them convex equally often. We can thus derive a Shape Behavior Model based on the observer’s own data (*SI Results*). Importantly, this Shape model no longer has constant performance across the four lighting directions (Fig. 3C).

For all cortical areas, Shape classification accuracy was above chance (0.5) for at least one of the four classifiers, or equivalently for at least one light direction. However, the performance of the classifiers was not uniform across all light directions (Fig. 3D). A repeated-measures ANOVA showed a significant main effect of light direction [$F(3,310) = 163.07, P < 0.0001$] and a significant interaction between ROI and light direction [$F(36,310) = 33.29, P < 0.0001$]. Looking at individual cortical areas, Shape classification accuracies depended on the light directions primarily in higher dorsal areas (dorsal visual and parietal regions) and not as much in ventral and early visual areas (Table S2).

It would be wrong to conclude that a certain brain region is making a shape discrimination simply from the fact that our Shape classifier produces above-chance performance. As we saw in Fig. 3B, there is information in the pixel intensities to perform this task. The evidence that will support our claim that a brain region processes shape is the similarity of the pattern of performance for the four lighting directions (Fig. 3D) with that of our Shape Behavior Model (Fig. 3C). This analysis is presented in a later section.

Light classification. Using similar analysis methods (MVPA) as above, we investigated which cortical areas contain information that allows us to discriminate light position (left vs. right) from fMRI signals. Four Light classifiers were trained to discriminate whether the viewed image was lit from the left or the right (Fig. 4A). The first classifier performed a classification between left and right light directions when the light was at 22.5° from the vertical and the object was convex (i.e., images “b” and “g” in Fig. 1A). The second classifier performed the same classification when the light was directed 67.5° from the vertical and the object was convex. The remaining two classifiers performed the same classification for concave objects (Fig. 4A).

Similar to the analysis for the Shape classifiers, we tested the performance of a model based on the image intensities (Fig. 4B) and another model based on the behavioral data (Fig. 4C). The details of these models are given in the *SI Results*. It is important to note that both models’ performance varied with light direction, with best performance for light directions near the horizontal and worst performance for light directions near the vertical. Critically, the performance across lighting directions formed a symmetrical U-shape pattern for the image-based model and an asymmetrical pattern for the behavioral model. The asymmetry of the Light behavior model results from the left bias of the estimated light direction (Fig. 1C).

The Light classification accuracies were high primarily in retinotopic areas, and they varied the most across the four classifiers

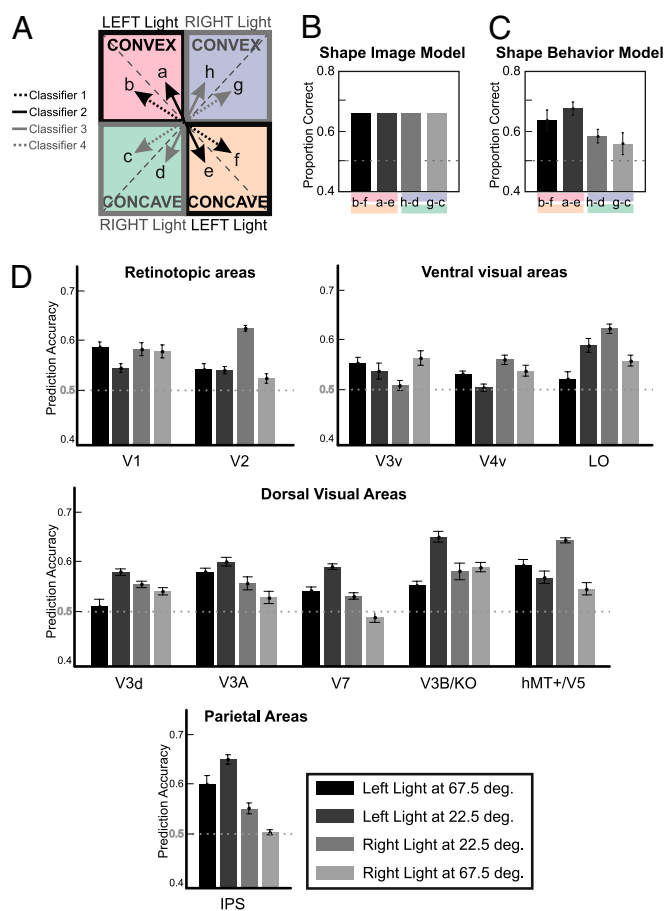


Fig. 3. Shape Classifier. MVPA for the classification of Shape (convex vs. concave) from fMRI data. (A) Classifier 1 compares activations for convex and concave stimuli when light is located 67.5° on the left: that is images “b” and “f” shown in Fig. 1A. Classifier 2 compares activations for stimuli when the light was 22.5° on the left: that is for images “a” and “e,” classifier 3 for images “h” and “d,” and classifier 4 for images “g” and “c” (Fig. 1). (B) Expected classification accuracies of the Shape Image Model. This model takes into account only the variation in pixel intensities between a pair of images. (C) Expected classification accuracies of the Shape Behavior Model. This model is based on the ability to discriminate convex and concave shapes as measured behaviorally for each observer (Fig. 1C) and predicts a nonuniform performance across the four lighting conditions. (D) Classification accuracies for the four lighting directions across ROI. The mean classification accuracy is based on 100 voxels per area. Error bars indicate SEM across observers ($n = 7$). The dashed line indicates the chance classification level (50%).

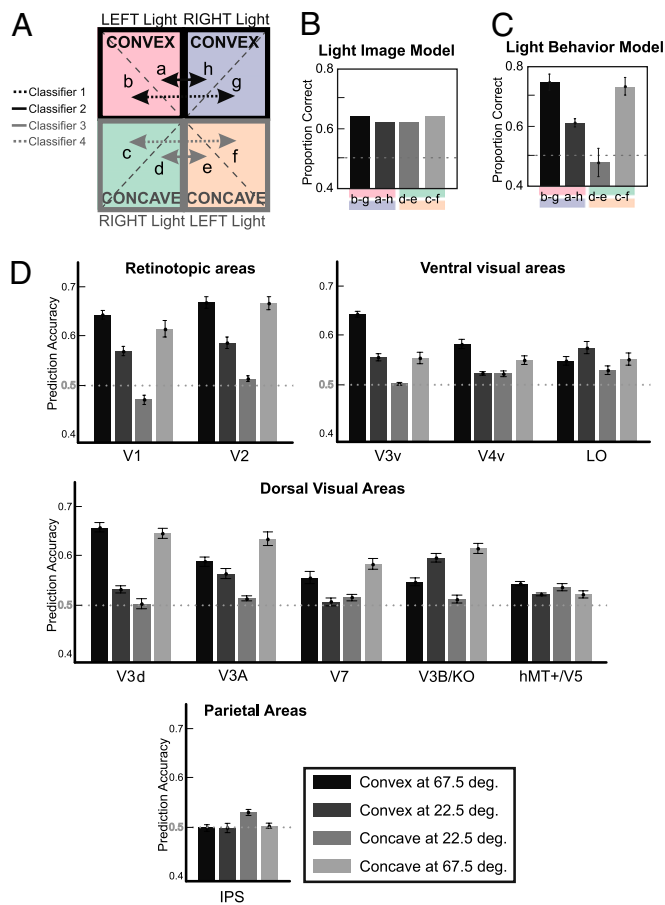


Fig. 4. Light Classifier. MVPA for the classification of Light (left vs. right). (A) Classifier 1 compares activations related to left-lit convex stimuli (67.5°) vs. right-lit convex stimuli (67.5°). Classifier 2 compares activations related to left-lit convex stimuli (22.5°) vs. right-lit convex stimuli (22.5°). Classifier 3 compares activations related to right-lit concave stimuli (22.5°) vs. left-lit concave stimuli (22.5°). Finally, Classifier 4 compares activations related to right-lit concave stimuli (67.5°) vs. left-lit concave stimuli (67.5°). (B) Expected classification of the Light Image Model. This model takes only into account the variation in pixel intensities between a pair of images. (C) Expected classification accuracies of the Light Behavior Model. This model performs a Bayesian inference based on the prior assumption that light comes from above-left as measured behaviorally for each observer (Fig. 1C) and predicts an asymmetrical performance across the four lighting conditions. (D) Classification accuracies for each of the four lighting directions. Mean classification accuracy is based on 100 voxels per area. Error bars indicate SEM across observers ($n = 7$). The dashed line indicates the chance classification level (50%).

also in these early visual areas (Fig. 4D). A repeated-measures ANOVA showed a significant interaction between ROI and classifier number where the classifier numbers are illustrated in Fig. 4A [$F(36,312) = 33.86, P < 0.0001$]. Looking at individual cortical areas, Light classification accuracies depended on light directions primarily in early visual areas and less so in parietal and dorsal visual areas (Table S2).

Relationship between behavior-based models and neural processing for Shape vs. Light. The perception of shape-from-shading involves an appropriate correspondence between the dark and bright image contours and the light source position, and an appropriate grouping of the contours to form a convex or concave shape. The observer's ability to complete the first stage of processing is well characterized by the Light model we have developed above (Fig. 4C), whereas the ability to complete the second stage is characterized by the Shape model (Fig. 3C). Obviously an observer is

able to complete both stages of processing, as illustrated by the fact that both models are based on the same behavioral data (Fig. 1C). However, some brain areas might be involved more in one of the two stages of processing.

To characterize the extent to which a specific brain area is involved in processing Shape or Light, we calculated the Spearman rank correlation of classification accuracies with the corresponding Shape and Light behavior models. Fig. 5A illustrates these correlations for V1 and regions along the IPS. For example, each symbol in the *Upper Left* plot shows the performance of the four light classifiers for one observer with the Light model for that observer. Even though the figure shows all observers together, correlations were computed individually for each observer. Across all observers, the Light classifier accuracies were well predicted by the Light model, as illustrated by a strong correlation (Spearman's rho: $\rho = 0.46$). This finding suggests that V1 is involved in the processing of light orientation. However, a similar analysis showed no significant correlation of the Shape classifier accuracies with the Shape model ($\rho = 0.04$), suggesting that V1 is not involved in 3D shape processing from shading. The *Lower* row of Fig. 5A illustrates the opposite effect for areas along the IPS, namely

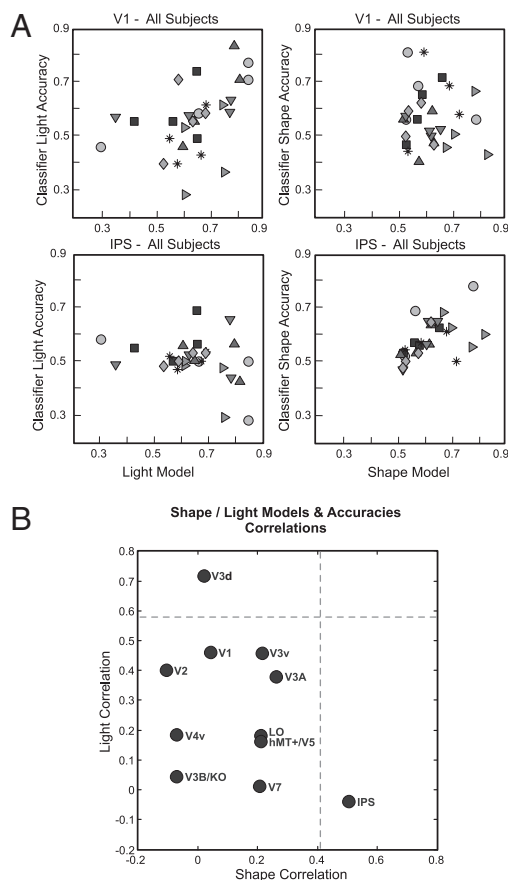


Fig. 5. Correlating pattern classification and behavior-based models. (A) Correlations of classification accuracies (Light classifier on the *Left* and Shape classifier on the *Right*) and the respective behavioral models for V1 (*Upper*) and regions along the IPS (*Lower*). Classification accuracies for V1 correlate well with the Light model but not with the Shape model, whereas accuracies for IPS correlate with Shape but not with Light. (B) Summary plot of correlations between classifier accuracies and Shape and Light models for each ROI. For each cortical area, the Spearman correlation of the Light classifiers with the Light Behavior Model is plotted against the correlation of the Shape classifiers with the Shape Behavior Model. Dashed lines represent significance criteria for $P = 0.05$ based on a permutation analysis constrained to the image information.

a strong involvement in Shape processing ($\rho = 0.51$) but not in Light processing ($\rho = -0.04$).

Fig. 5*B* illustrates the correlations of the Shape and Light classifier accuracies across cortical areas with the respective behavior models. The right-hand side of the figure contains areas that are involved in Shape processing, whereas the top part of the figure contains areas involved in Light processing. The dashed lines represent correlation criteria to reach a significance P value of 0.05 according to a permutation analysis constrained to the Shape and Light image models (30). In this test, all permutations were allowed for the Shape analysis, but only permutations within the group composed of the classifiers “b–g” and “c–f” on the one hand and within the group “a–h” and “d–e” on the other hand were allowed for the Light analysis. Early retinotopic areas, in particular V1, V2, and V3d appear to be involved in Light processing, although only the correlation of V3d reached significance ($P = 0.048$). In contrast, higher areas appear to be involved in Shape processing: V3v, LO, V3A, V7, hMT+/V5, and regions along the IPS, although only the correlation of IPS reached significance in our permutation test ($P < 0.01$). The result that areas traditionally involved in object processing (such as LO) did not reach significance may be because of the larger complexity of the objects used in these previous studies (e.g., ref. 16). Critically, we did not find any cortical area that showed a significant correlation with both the Light and Shape models, suggesting that no area was fully involved in both Light and Shape processing. Dissociating feed-forward fMRI signals that may relate to light processing from activity related to shape perception in higher areas is complicated due to the complex nature of the blood-oxygen level-dependent (BOLD) response. However, we found a surprisingly clean division of labor between early visual areas that process image-contrast properties and higher visual areas that correlate with the perceived object shape.

Discussion

Our study investigated the neural network that allows us to infer 3D shape from shading. In this inference, humans assume that light is coming from above their head, but also slightly from the left (1, 2). Our findings reveal that this prior knowledge on illumination is processed early in the visual system, primarily in early retinotopic areas. We reached this conclusion by capitalizing on the left bias for the assumed light-source position. Because of this left bias, flipping an image about the vertical axis led to changes in the reliability of the perceived 3D shape depicted by the shading cue. Thus the same shape, for example a convex ring, was perceived more or less frequently depending on whether light was simulated on the left or the right. These variations in the image interpretation depend neither on the image itself (an image and its mirror reverse contain the same information) nor on the shape (similar variations occur for a concave ring) but purely on the characteristics of the assumed light-source position (a left bias). Activity in early visual areas was not only affected by lighting direction (left or right), but it was modulated in a manner very similar to the human behavior. In contrast, the activity in higher visual areas (in particular in parietal areas) was not affected by lighting direction (left vs. right), but instead reflected the perceived shape (convex vs. concave). Past studies had already highlighted several regions of the shape-from-shading network, in particular early visual areas (5–9) and to a lesser extent higher extrastriate (10) and parietal areas (11). Our study helps delineating the functions mediated by these different cortical areas. To obtain these results, we have used computational modeling to provide a strong link of fMRI BOLD signals with behavior.

Are the effects we observe in early visual areas the signature of the representation of the assumed light-source position, or is this prior knowledge represented somewhere else and fed back to early visual areas? Intuitively, it is reasonable to expect that some prior knowledge, especially if it can be verbalized (i.e., “light is

coming from above because the sun is above our head”), should be represented in higher cortical areas where semantic information can be stored. According to this popular view, prior knowledge is processed in a top-down manner, following for instance the elegant Bayesian-belief propagation model where “the feedforward input drives the generation of the hypotheses, and the feedback from higher inference areas provides the priors to shape the inference at the earlier levels” (21). However, our results are inconsistent with the feedback hypothesis. Should the prior knowledge be processed in higher visual areas, both these areas and the early visual areas should contain information sensitive to the consequences of changing the light direction from left to right. Because this sensitivity was lacking in higher visual areas, we conclude that the prior knowledge on illumination is processed in the early visual areas. This finding is an important result exemplifying that prior knowledge does not necessarily imply high-level representation and top-down processing as previously suggested. More specifically, the prior knowledge about the world that guides our perception of complex 3D environments, such as the assumption that light is coming from above and slightly from the left, appears to be one of these constraints represented early in the visual system. These results are consistent with previous studies that have shown a short delay for the neural structures involved in shape-from-shading (7). The reason for the left, rather than right, bias is still unresolved, but some speculations have been proposed that would be consistent with an early processing of this prior knowledge (31).

We propose that the representation of prior knowledge about the illumination direction early in the visual system is important, as light affects the contrast of an image (32); that is, a basic visual property that needs to be processed before more interesting image understanding takes place. Other types of visual prior knowledge do not necessarily have to be represented in early visual areas. More generally, we expect that prior knowledge will be represented wherever the property related to the prior is explicitly processed. For example, preferences for slow motion (33) might also be represented in early visual areas but preferences for looming should be searched in areas sensitive to this more complex kind of motion [i.e., the human homolog of area medial superior temporal in macaque monkeys (34)]. Similarly, preferences for convex rather than concave objects (e.g., ref. 35) should be searched in areas involved in shape perception, in particular in areas hMT+/V5 (36), LO, and IPS, as suggested by our findings. Taking into account the role of prior knowledge in our perceptual decisions will advance our understanding of the computations underlying human perception across cortical areas.

Methods

Psychophysics. Before scanning, we conducted a psychophysical study to confirm that the stimuli were perceived as 3D by the observers when lying in the scanner. Behavioral data were collected in a mock scanner with the same stimulus delivery equipment as in the scanner. Each session [composed of all 96 stimuli: two shapes (convex or concave) \times four light directions ($\pm 22.5^\circ$ and $\pm 67.5^\circ$) \times six sector positions \times two levels of blur] was repeated six times. All stimuli were presented in a randomized sequence. Stimuli were shown for 100 ms and then followed immediately by a mask. No feedback was provided to the observers. All subjects were instructed to detect the shape (convex or concave) of the odd sector. Specific details on the observers and stimuli can be found in *SI Methods*.

fMRI Design and Procedure. Cortical regions of interest were identified by presenting observers with 48 ring stimuli and 24 scrambled versions of these stimuli. Data used in the MVPA were collected on separate runs where only the ring stimuli were presented. Observers were instructed to perform a detection task on the fixation (i.e., detect a change from “+” to “x”) and performed with an accuracy larger than 80%. Each stimulus was presented for 200 ms followed by a blank interval (600 ms).

fMRI Data Acquisition. All experiments were conducted using a 3-Tesla Philips Achieva MRI scanner at the Birmingham University Imaging Centre. T2*-weighted

functional and T1-weighted anatomical ($1 \times 1 \times 1$ mm) data were collected with an eight-channel SENSE head coil. EPI data were acquired from 33 axial slices (whole-head coverage, TR: 2,000 ms, TE: 35 ms, flip-angle: 80°, $2.5 \times 2.5 \times 3$ mm resolution).

fMRI Data Analysis. The fMRI data were analyzed using Brain Voyager QX (Brain Innovations). Preprocessing of the functional data consisted of slice-scan time correction, head movement correction, temporal high-pass filtering (three cycles), and linear trend removal. No spatial smoothing was applied except in the group data analysis (Gaussian filter; full-width at half maximum, 6 mm) that was conducted for identifying cortical areas involved in the processing of shape from shading. Anatomical data were used for 3D cortex reconstruction, inflation, and flattening. Functional images were aligned to anatomical data and the complete data transformed into Talairach space. **MVPA.** For the MVPA (e.g., refs. 22, 23), we used linear SVM classifiers and followed cross-validation procedures as in our previous studies (37). For each ROI (retinotopic areas, V3B/KO, and hMT+/V5, shape-from-shading responsive regions), we sorted the voxels according to their response (t -statistic) to all stimulus conditions compared with fixation baseline across all experimental runs. We selected the same number of voxels across ROIs and observers by restricting the pattern size to those voxels that showed a significant ($P < 0.05$) t value when comparing all stimulus conditions vs. fixation. This procedure resulted in the selection of 100 voxels per ROI, which is comparable to the dimensionality used in previous studies (22). The mini-

mum number of voxels included for each ROI and subject was 100. We normalized (z -score) each voxel time-course separately for each experimental run to minimize baseline differences across runs. The data vectors for the multivariate analysis were generated by shifting the fMRI time series by 4 s to account for the delay of the hemodynamic response and then averaging all time series datapoints of one experimental block. We selected data vectors according to the comparison of interest (Figs. 3A and 4A) and split them into a training sample comprising the data of seven runs and a test sample comprising the remaining run (for one of the observers, we collected six experimental runs: five runs were used as training sample and one run as test). We performed an 8-fold cross-validation leaving one run out (test sample). For each subject, we averaged the accuracy rates (number of correctly assigned test patterns/total number of assignments) across the cross-validation runs. We evaluated the statistical significance across subjects using repeated measures ANOVA (Figs. S2–S4 and Tables S3 and S4).

ACKNOWLEDGMENTS. We thank Sheng Li and Dirk Ostwald for helpful suggestions and discussions on the analysis, Matthew Dexter for help with the eye movement analysis software, and Ken Knoblauch and Andrew Welchman for comments on an earlier draft. This work was supported by an International Short Visit Royal Society Fellowship (2006) (to Z.K. and P.G.), a Chaire d'Excellence from the French Ministry of Research (to P.M.), a European Commission FP7 grant PITN-GA-2008-214728 (CODDE) (to P.M. and Z.K.), and Cognitive Foresight Initiative Grant BB/E027436/1 (to Z.K.).

1. Sun J, Perona P (1998) Where is the sun? *Nat Neurosci* 1:183–184.
2. Mamassian P, Goutcher R (2001) Prior knowledge on the illumination position. *Cognition* 81:B1–B9.
3. Ramachandran VS (1988) Perception of shape from shading. *Nature* 331:163–166.
4. Mamassian P (2008) Ambiguities and conventions in the perception of visual art. *Vision Res* 48:2143–2153.
5. Humphrey GK, et al. (1997) Differences in perceived shape from shading correlate with activity in early visual areas. *Curr Biol* 7:144–147.
6. Lee T-S, Yang CF, Romero RD, Mumford D (2002) Neural activity in early visual cortex reflects behavioral experience and higher-order perceptual saliency. *Nat Neurosci* 5:589–597.
7. Mamassian P, Jentzsch I, Bacon BA, Schweinberger SR (2003) Neural correlates of shape from shading. *Neuroreport* 14:971–975.
8. Hou C, Pettet MW, Vildavski VY, Norcia AM (2006) Neural correlates of shape-from-shading. *Vision Res* 46:1080–1090.
9. Smith MA, Kelly RC, Lee TS (2007) Dynamics of response to perceptual pop-out stimuli in macaque V1. *J Neurophysiol* 98:3436–3449.
10. Georgieva SS, Todd JT, Peeters R, Orban GA (2008) The extraction of 3D shape from texture and shading in the human brain. *Cereb Cortex* 18:2416–2438.
11. Taira M, Nose I, Inoue K, Tsutsui K-i (2001) Cortical areas related to attention to 3D surface structures based on shading: an fMRI study. *Neuroimage* 14:959–966.
12. Cox DD, Savoy RL (2003) Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* 19:261–270.
13. Haynes J-D, Rees G (2006) Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7:523–534.
14. Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
15. Backus BT, Fleet DJ, Parker AJ, Heeger DJ (2001) Human cortical activity correlates with stereoscopic depth perception. *J Neurophysiol* 86:2054–2068.
16. Kourtzi Z, Erb M, Grodd W, Bülthoff HH (2003) Representation of the perceived 3-D object shape in the human lateral occipital complex. *Cereb Cortex* 13:911–920.
17. Welchman AE, Deubelius A, Conrad V, Bülthoff HH, Kourtzi Z (2005) 3D shape perception from combined depth cues in human visual cortex. *Nat Neurosci* 8:820–827.
18. Tyler CW, Likova LT, Kontsevich LL, Wade AR (2006) The specificity of cortical region KO to depth structure. *Neuroimage* 30:228–238.
19. Preston TJ, Kourtzi Z, Welchman AE (2009) Adaptive estimation of three-dimensional structure in the human brain. *J Neurosci* 29:1688–1698.
20. Gerardin P, de Montalembert M, Mamassian P (2007) Shape from shading: New perspective from the Polo Mint stimulus. *J Vis* 7(11):13, 1–11.
21. Lee T-S, Mumford D (2003) Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis* 20:1434–1448.
22. Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
23. Haynes J-D, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
24. Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
25. Preston TJ, Li S, Kourtzi Z, Welchman AE (2008) Multivoxel pattern selectivity for perceptually relevant binocular disparities in the human brain. *J Neurosci* 28:11315–11327.
26. Haxby JV, et al. (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430.
27. Hanson SJ, Matsuka T, Haxby JV (2004) Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a “face” area? *Neuroimage* 23:156–166.
28. O’Toole AJ, Jiang F, Abdi H, Haxby JV (2005) Partially distributed representations of objects and faces in ventral temporal cortex. *J Cogn Neurosci* 17:580–590.
29. Williams MA, Dang S, Kanwisher NG (2007) Only some spatial patterns of fMRI response are read out in task performance. *Nat Neurosci* 10:685–686.
30. Pesarin F (2001) *Multivariate Permutation Tests with Applications to Biostatistics* (John Wiley & Sons, Chichester).
31. Mamassian P (2004) Impossible shadows and the shadow correspondence problem. *Perception* 33:1279–1290.
32. Pereverzeva M, Murray SO (2008) Neural activity in human V1 correlates with dynamic lightness induction. *J Vis* 8(15):8, 1–10.
33. Stocker AA, Simoncelli EP (2006) Noise characteristics and prior expectations in human visual speed perception. *Nat Neurosci* 9:578–585.
34. Wall MB, Smith AT (2008) The representation of egomotion in the human brain. *Curr Biol* 18:191–194.
35. Mamassian P, Landy MS (1998) Observer biases in the 3D interpretation of line drawings. *Vision Res* 38:2817–2832.
36. Kourtzi Z, Bülthoff HH, Erb M, Grodd W (2002) Object-selective responses in the human motion area MT/MST. *Nat Neurosci* 5:17–18.
37. Li S, Ostwald D, Giese M, Kourtzi Z (2007) Flexible coding for categorical decisions in the human brain. *J Neurosci* 27:12321–12330.