

Representation of Maximally Regular Textures in Human Visual Cortex

Peter J. Kohler,¹ Alasdair Clarke,² Alexandra Yakovleva,¹ Yanxi Liu,³ and Anthony M. Norcia¹

¹Department of Psychology, Stanford University, Stanford, California 94305, ²School of Psychology, University of Aberdeen, Aberdeen AB24 2UB, United Kingdom, and ³Departments of Computer Science and Engineering, and Electrical Engineering, Pennsylvania State University, University Park, Pennsylvania 16802

Naturalistic textures with an intermediate degree of statistical regularity can capture key structural features of natural images (Freeman and Simoncelli, 2011). V2 and later visual areas are sensitive to these features, while primary visual cortex is not (Freeman et al., 2013). Here we expand on this work by investigating a class of textures that have maximal formal regularity, the 17 crystallographic wallpaper groups (Fedorov, 1891). We used texture stimuli from four of the groups that differ in the maximum order of rotation symmetry they contain, and measured neural responses in human participants using functional MRI and high-density EEG. We found that cortical area V3 has a parametric representation of the rotation symmetries in the textures that is not present in either V1 or V2, the first discovery of a stimulus property that differentiates processing in V3 from that of lower-level areas. Parametric responses were also seen in higher-order ventral stream areas V4, VO1, and lateral occipital complex (LOC), but not in dorsal stream areas. The parametric response pattern was replicated in the EEG data, and source localization indicated that responses in V3 and V4 lead responses in LOC, which is consistent with a feedforward mechanism. Finally, we presented our stimuli to four well developed feedforward models and found that none of them were able to account for our results. Our results highlight structural regularity as an important stimulus dimension for distinguishing the early stages of visual processing, and suggest a previously unrecognized role for V3 in the visual form-processing hierarchy.

Key words: hierarchical processing; mid-level vision; perceptual organization; symmetry; texture; ventral stream

Significance Statement

Hierarchical processing is a fundamental organizing principle in visual neuroscience, with each successive processing stage being sensitive to increasingly complex stimulus properties. Here, we probe the encoding hierarchy in human visual cortex using a class of visual textures—wallpaper patterns—that are maximally regular. Through a combination of fMRI and EEG source imaging, we find specific responses to texture regularity that depend parametrically on the maximum order of rotation symmetry in the textures. These parametric responses are seen in several areas of the ventral visual processing stream, as well as in area V3, but not in V1 or V2. This is the first demonstration of a stimulus property that differentiates processing in V3 from that of lower-level visual areas.

Introduction

Objects and surfaces in the visual world are defined not just by their shapes and orientation, but also by their texture and surface markings. Textures are characterized by the repetitive occurrence of a pattern that may appear more or less regular at different

scales (Arcizet et al., 2008). Textures contribute strongly to the perception of the material properties of objects (Adelson, 2001), and to the perception of both patterns (Kass and Witkin, 1987; Dakin and Watt, 1997) and 3D shape (Knull, 1998a,b,c; Li and Zaidi, 2000).

Image texture can be understood in the context of regularity space, which positions any texture on a continuum going from low to high regularity (Liu et al., 2004b). “Microtextures”—image patches whose appearance is invariant with phase information (Galerne et al., 2011)—are found at the low-regularity end of the spectrum. New, perceptually indistinguishable, microtexture patches can be generated by simply randomizing the phase of the original patch while retaining the amplitude spectrum. This class of natural image texture includes gravel, sand, clouds, and waves, but other natural textures have more struc-

Received Aug. 3, 2015; revised Nov. 16, 2015; accepted Nov. 23, 2015.

Author contributions: P.J.K., A.Y., Y.L., and A.M.N. designed research; P.J.K. and A.C. performed research; P.J.K. analyzed data; P.J.K., A.C., A.Y., Y.L., and A.M.N. wrote the paper.

This research was supported by National Science Foundation INSPIRE Grant 1248076, which was awarded to Y.L. and A.M.N.

The authors declare no competing financial interests.

Correspondence should be addressed to Peter J. Kohler, Department of Psychology, Stanford University, Jordan Hall, Building 420, 450 Serra Mall, Stanford, CA 94305. E-mail: pjkohl3r@gmail.com.

DOI:10.1523/JNEUROSCI.2962-15.2016

Copyright © 2016 the authors 0270-6474/16/360714-16\$15.00/0

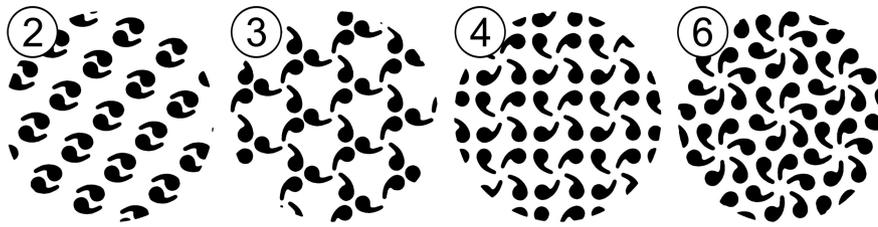


Figure 1. Tutorial examples of the four distinct wallpaper patterns used in the experiment, containing no reflection or glide symmetries. The maximum order of rotation symmetry for each wallpaper group is indicated next to each image. In these examples, the fundamental region is a comma-like symbol (illustration adapted from Wade, 1993).

tured, repeating elements. These textures cannot be synthesized on the basis of the power spectrum alone and have been called “macrot textures” to highlight the increased regularity over micro-textures (Galerne et al., 2011).

An influential approach to synthesizing macrot textures has been to use a set of joint statistics over a set of filters tuned for orientation, location, and scale (Portilla and Simoncelli, 2000). This biologically inspired synthesis approach to texture has been applied as an analysis tool for understanding texture representations in visual areas V2 (Freeman et al., 2013) and V4 (Okazawa et al., 2015). While neurons in V1 are largely insensitive to the information captured by the synthesis model, sensitivity to the higher-order features is present starting in V2 in both the macaque single-unit data and in human fMRI (Freeman and Simoncelli, 2011; Freeman et al., 2013).

Here, we investigated a class of textures that are found at the extreme end of the texture regularity spectrum. These maximally regular textures each belong to 1 of the 17 crystallographic “wallpaper groups,” which are defined by a unique combination of the following four fundamental symmetries: translations, rotations, reflections (i.e., mirror symmetry), and glide reflections (Fedorov, 1891; Polya, 1924; Liu et al., 2010). All groups contain translation symmetry, which arises from a group-specific textural motif that periodically tiles the plane. In an effort to determine how and where these higher-order regularities are represented in human visual cortex, we focused on a subset of this class of textures that all contain rotation as well as translation, but no other symmetries. Previous imaging work on symmetry has identified several extrastriate areas that are sensitive to reflection symmetry in dot patterns (Sasaki et al., 2005; Tyler et al., 2005). A subset of these regions remained sensitive even when attention was controlled, while others had a significant response during only passive viewing (Sasaki et al., 2005).

Here, we extended this work using a combination of functional MRI and EEG source imaging, and found that visual areas V3, V4, VO1, and the lateral occipital complex (LOC) each had a parametric response to rotation symmetry. This sensitivity to rotation symmetry was absent in V1 or V2, suggesting that the information conveyed by this form of regularity is not encoded before V3. To investigate whether the symmetry responses in V3 and V4 rely on top-down signals from object areas later in the ventral stream, we performed an EEG source localization experiment that allowed us to determine the temporal order of symmetry processing across visual areas. These measurements indicated that the onset of the symmetry response occurs earlier in V3 and V4 than in LOC, suggesting that symmetry information is propagated to the ventral stream in a bottom-up fashion. Our results show that regular textures are represented parametrically starting in V3 and throughout ventral temporal cortex.

Materials and Methods

Participants. Twenty-five participants (11 females; mean age, 28.7 ± 13.3) took part in the main EEG experiment. Twelve of these participants (6 females; mean age, 34.8 ± 14.0) also took part in the fMRI experiment. Ten additional participants (4 females; mean age, 25.3 ± 9.0) took part in an additional control EEG experiment, with 3 people participating in all three experiments. All were prescreened to confirm that they had normal or corrected-to-normal visual acuity on the Bailey-Lovie chart and normal stereopsis on the RandDot test (<http://precision-vision.com/products/stereo-vision-tests/randdot-stereo-test.html>).

Their informed consent was obtained before the experiment under a protocol that was approved by the Institutional Review Board of Stanford University.

Wallpaper stimulus generation. Wallpaper patterns are repetitive 2D patterns that tile the plane. There are 17 unique wallpaper patterns, corresponding to the set of planar symmetry groups (Fedorov, 1891; Polya, 1924; Liu et al., 2010). Each of the 17 wallpaper groups is built on one of five types of “unit lattices” that are used for tiling the plane without gaps. The “fundamental region” is the smallest repeating region in the wallpaper patterns. Within the unit lattice, multiple rigid transformations are applied to the fundamental region, which give rise to symmetries within the wallpaper group. As mentioned earlier, each wallpaper group contains a distinct combination of the following four fundamental symmetries: translations, rotations, reflections, and glide reflections.

In the experiments presented here, we measured neural responses to a subset of four of the 17 wallpaper groups: P2, P3, P4, and P6. These four groups all contain rotation symmetries, as well as translation symmetries given by the tiling of the unit lattice, but no other symmetries. Rotation symmetry around a particular point can be defined in terms of its order n , which means that the fundamental region can be rotated by an angle $360^\circ/n$ without changing. The four wallpaper groups used in the present experiments are illustrated schematically in Figure 1, using a comma-shaped symbol as the fundamental region. Each group contains rotation symmetry around several points that vary in order. For P2, the maximum order of rotation symmetry is 2; for P3, it is 3; for P4, it is 4; and for P6, it is 6.

We generated exemplars from the four wallpaper groups by using random noise textures as the fundamental regions based on a modification of the stimuli developed by Clarke et al. (2011b). In Figure 2, we show an exemplar from each of the four wallpaper groups, and indicate the fundamental region, rotation symmetry centers, and unit lattice on a magnified region of each texture. Although all four exemplars are generated from a fundamental region containing random noise, the resulting wallpapers are quite visually appealing, and very different from one another. Importantly, using random noise in the fundamental region means that for each wallpaper group an almost infinite number of distinct exemplars can be easily generated (a range of exemplars from group P6 are shown in Fig. 4).

Generating wallpaper patterns by directly tiling the unit lattice is algorithmically challenging because unit lattices can have non-right angles, making them difficult to concatenate on a raster graphics display. We circumvented this issue by generating the smallest rectangular tile that contains the unit lattice and that can be easily replicated, here referred to as the “repeating tile.” The generation of a repeating tile starts with the fundamental region, which is cut from a larger “base texture.” The base texture is generated by convolving spatial white noise with a Gaussian-like filter constrained by the dimensions (width and height) of the base texture.

The fundamental region forms the basis for the unit lattice, which in turn constrains the repeating tile. Because of this, the shape and size of the fundamental region depends on the wallpaper group being portrayed as well as the area of the repeating tile. The generation of the fundamental region for group P6, an isosceles triangle with angles 30° , 30° , and 120° , is

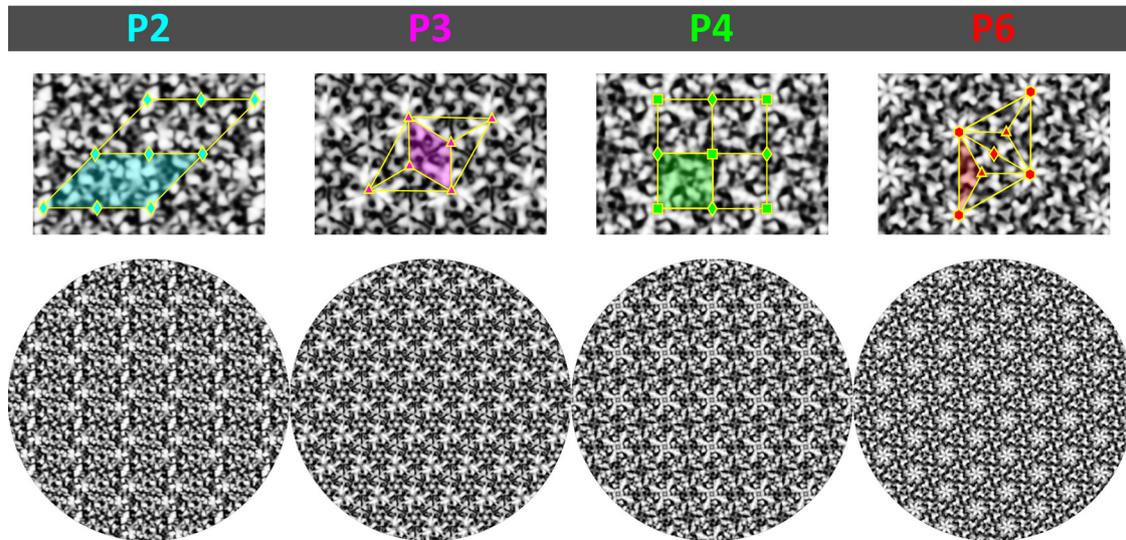


Figure 2. The stimuli used in the experiments reported here. Bottom, An exemplar from each of the four wallpaper groups. Top, Part of each exemplar, magnified so that the details of the pattern are easier to see. The fundamental regions are indicated with colored shading. The rotation symmetries in each pattern are indicated as follows: rotation of order 2 (rhombus), rotation of order 3 (triangle), rotation of order 4 (square), and rotation of order 6 (hexagon). The unit lattice is indicated with a yellow outline. Note how the fundamental region is rotated and repeated within the unit lattice, which is then used to tile the plane in each of the four exemplars. In the case of P3, the fundamental region is cut in half, and each half is rotated and repeated separately.

shown in Step 1 of Figure 3. After the fundamental region is generated, it is then rotated and repeated to generate the repeating tile appropriate for a specific wallpaper group. This process is illustrated in Steps 2–5 of Figure 3. The finished repeating tile is then used to tile the plane (Fig. 3, Step 6), before a final postprocessing procedure is applied (Fig. 3, Step 7), as described below.

The dimensions of the fundamental region define the area of the repeating tile, which in turn defines the spatial scale of the exemplar. A key issue when designing the wallpapers was to parameterize the dimensions of the fundamental region such that spatial scale was uniform across different groups. Because of forced relationships among the area of the unit lattice, the area of the fundamental region, and the dimensions of the repeating tile, it is not possible to equate all parameters across all groups. For our stimuli, we decided to equalize the area of the repeating tile across groups to 100×100 image pixels. This meant that the unit lattices had an area of 100^2 pixels for the P2 and P4 groups, and $2/3 \times 100^2$ for the P3 and P6 groups (Fig. 2). The initial horizontal and vertical dimensions of the base texture were set to the square root of the tile area for all groups to ensure that the granularity within the textures was consistent. We calculated the area of the repeating tile as a function of the dimensions of the fundamental region for each group and cropped the base texture so that the repeating tile would have the desired area. For the wallpaper groups containing rotations of order 3 and 6 (P3 and P6 groups), an additional upscaling operation was performed: after cropping, the base texture was scaled up multiple times (usually 9–10 times), allowing us to concatenate the results of rotation with high precision (within 3 pixels). The finished repeating tile was then scaled back down so that the area of the repeating tile was matched across all groups.

For each of the four wallpaper groups, we generated 20 exemplars, each starting with a new randomly generated base texture and resulting fundamental region. We then applied the following postprocessing procedure to each of the exemplars to accomplish several goals. First, we performed a Fourier decomposition on each of the 20 exemplar images and computed the average power spectrum across all exemplars within each group. We then replaced the power spectrum of each individual exemplar with the average, so that the power spectrum was equated across all exemplars. To minimize edge artifacts due to tiling, each exemplar image was then low-pass filtered using a Gaussian filter with size 9×9 pixels and SD $\sigma = 1$ pixel. Using the `histeq` function in MATLAB, we performed histogram equalization to increase the visual salience of the patterns, and then scaled the gray-level pixel intensities to 2–255. Finally, we applied a circular aperture mask to the images to minimize interac-

tions between the wallpaper group lattices and the pattern edges. The effect of the postprocessing procedure can be observed by comparing Steps 6 and 7 in Figure 3.

Control stimulus. Phase-randomized control exemplars were generated for each of the 20 exemplars from each wallpaper group. Control exemplars had the same power spectrum as the exemplar images for each group. Although the phase scrambling eliminates the symmetry relationships within the unit lattice, the periodicity of the lattice itself is inherited by the phase-scrambled images. This means that all control exemplars, regardless of which wallpaper group they are derived from, degenerate to another symmetry group, namely P1. P1 is the simplest of the wallpaper groups and contains only translations of a region whose shape derives from the lattice. Because the different wallpaper groups have different lattices, these P1 controls are slightly different between groups. Among the four groups we studied here, P2 and P4 shared one lattice, and P3 and P6 shared another, different lattice. As we shall see, our experimental design takes these lattice differences into account by comparing the neural responses evoked by a wallpaper group to that evoked by the matched control exemplars that share the same lattice structure, but have no rotation symmetry.

Stimulus presentation. In the functional MRI experiment, the wallpaper images were shown on a 47 inch Resonance Technology LCD display. Participants viewed the screen through a mirror at a distance of 277 cm. The screen had a resolution of 1024×768 pixels, an 8 bit color depth, and a refresh rate of 60 Hz. The mean luminance was 34.39 cd/m^2 , and contrast was 90%. The circular aperture defining the edge of the wallpapers had a diameter of 11.2° of visual angle, and the rest of the screen was uniformly gray. In the EEG experiment, the stimuli were shown on a 24.5 inch Sony Trimaster EL PVM-2541 organic light-emitting diode display, with a screen resolution of 1920×1080 pixels, 8-bit color depth and a refresh rate of 60 Hz, viewed at a distance of 70 cm. The mean luminance was 69.93 cd/m^2 , and contrast was 95%. The diameter of the circular aperture was slightly larger in this experiment at 13.8° of visual angle, and again the background was gray.

Structural and functional MRI acquisition. Functional and structural MRI data were collected on a Discovery 750 MRI system (General Electric Healthcare) equipped with a 32-channel head coil (Nova Medical) at the Center for Cognitive and Neurobiological Imaging at Stanford University. For each participant, we acquired two whole-brain T1-weighted structural datasets ($1.0 \times 1.0 \times 1.0 \text{ mm}$ resolution, TE = 2.5 ms, TR = 6.6 ms, flip angle = 12° , FOV = $256 \times 256 \text{ mm}$) that were used for tissue segmentation and registration with the functional scans. We also acquired a single whole-brain, T2-weighted structural dataset ($1.0 \times 1.0 \times$

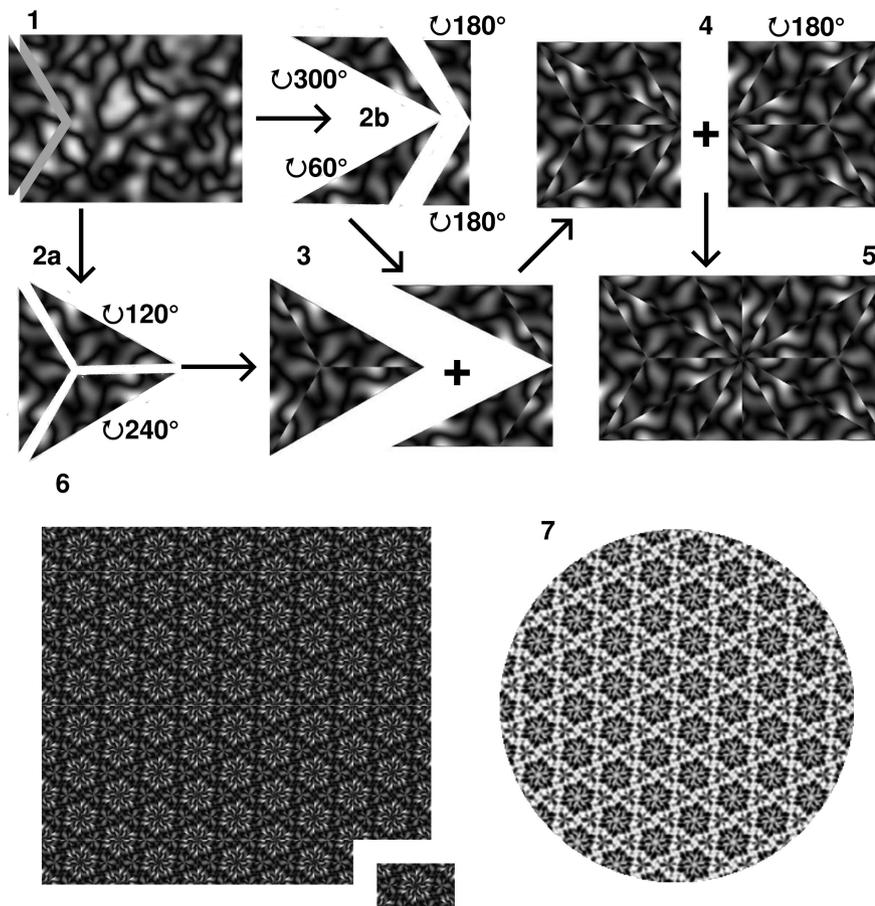


Figure 3. Diagram of stimulus generation steps for wallpaper group P6, as follows: 1, the fundamental region is selected from the filtered random-noise base texture; 2a, the fundamental region is rotated 120° and 240°, and combined with itself into an equilateral triangle; 2b, building the triangle “liner.” The fundamental region is rotated 180°; the top half is stacked at the bottom, and the bottom half goes to the top. This set is now combined with 300° and 60° rotated versions of the fundamental region. 3, Triangle and triangle liner are joined together to produce half of the repeating tile; 4, half of the repeating tile is rotated 180° and combined with itself to produce a rectangular repeating tile; 5, the finished repeating tile; 6, the repeating tile is used to tile the plane, repeated eight times down and five times across; and 7, the finished stimulus after postprocessing (power spectrum replacement, low-pass filtering, histogram equalization, and application of a circular aperture mask).

1.0 mm resolution, TE = 75 ms, TR = 2500 ms, flip angle = 90°, FOV = 256 × 256 mm), which was used for defining the tissue boundaries used for the EEG head model. For the functional MRI experiment, we collected functional data consisting of 30 coronal slices (2.0 × 2.0 × 2.0 mm resolution, TE = 30 ms, TR = 2000 ms, flip angle = 77°, FOV = 220 × 220 mm), positioned at an oblique angle to maximize the coverage of occipital, ventral, and parietal cortices. The scans used for localization of functionally defined regions of interest (ROIs) had the same voxel size and TE/TR as that used in the main experiment, but a multiplexed EPI sequence was used that allowed for whole-brain coverage.

fMRI experimental procedure. We used a block design for the fMRI experiment in which 12 s experiment blocks (referred to as “A blocks”) alternated periodically with 12 s control blocks (referred to as “B blocks”), yielding a 24 s base period for the paradigm, which was repeated 10 times in what we refer to as a “scan.” The paradigm is illustrated schematically in Figure 4. Ten stimulus cycles were shown per scan, with an additional half-cycle (one 12 s control block) being shown in the beginning of the scan to allow the brain and the scanner to settle. The data collected during this “dummy” period were not included in the analysis. Fourier analysis was used to quantify BOLD activation at frequencies equal to exact integer multiples of the 240 s scan period (0.0042 Hz) up to 0.5 Hz, as limited by the 2 s TR. The paradigm frequency was 10 cycles per scan, and all plots use this convention for the frequency axis.

During the experiment blocks, 10 wallpaper group exemplars were shown interleaved with their matched control exemplars, alternating at 1

Hz (e.g., 500 ms duration P1 control images followed by 500 ms presentation of the corresponding intact exemplar). During the control blocks, P1 control exemplars were alternated with another set of P1 control exemplars generated using the same procedure (for a total of 20 control images per block, with each image again presented for 500 ms). Experiment and control blocks differed only in the presence of intact rotation symmetry in the experiment block. Ten image pairs (control–exemplar or control–control) were shown in each block, always in the same order, and the first and last pairs were shown twice, for a total of 24 image presentations within a block. Within each scan, exemplars and control stimuli all came from the same wallpaper group, and we collected three 4 min scans for each group. To avoid the effects of scan order within a session, we scanned all four groups in a specific order, which was repeated three times. We always used the same repeating order (P2 → P3 → P4 → P6), but different starting groups were used for each participant, such that all four groups were used as the starting group for equal amounts of time. The fMRI experiment design is illustrated in Figure 4B.

We used a concurrent task to ensure that participants were alert and paying attention to the visual stimulus. Two times per block, for both A and B blocks, a randomly chosen image pair was shown at reduced contrast, and participants were instructed to press a button whenever they detected a contrast reduction. Before each scan, we adjusted the contrast reduction such that the average value for participant accuracy was kept at ~85% correct.

EEG experimental procedure. The EEG experiment was designed to mimic the fMRI experiment as closely as possible, while taking into account the superior temporal resolution of EEG. We used a steady-state design, in which P1 control images alternated with test images from each of the four wallpaper groups under study (P2, P3, P4, and P6). As in the fMRI study, each exemplar image was always preceded by its matched P1 control image, which was generated as described above. Each image was shown for 600 ms, and a control image followed by an exemplar image constituted a single stimulus cycle, with a frequency of 0.83 Hz. A trial consisted of 10 stimulus cycles corresponding to 10 different exemplar images and matched controls, which all came from the same wallpaper group. The images were always shown in the same order within a trial.

Participants initiated each trial with a button press, which allowed them to take breaks between trials. Trials from a single wallpaper group were presented in blocks of four repetitions, which were themselves repeated twice per session, and were shown in random order within each session. During each trial, participants performed the same concurrent task as in the fMRI experiment, as follows: two times per trial, an image pair was shown at reduced contrast, and the participants were instructed to press a button on the joypad. As in the fMRI experiment, we adjusted the contrast reduction such that the average value for participant accuracy was kept at ~85% correct.

fMRI analysis. After removing the dummy TR values, the fMRI data were preprocessed in AFNI (Cox, 1996). Preprocessing included the following steps: slice time correction; motion registration (the third TR of the first scan was always used as the base); scaling (each voxel was scaled to a mean of 100, and values were clipped at 200); and detrending [removing components corresponding to the six motion registration parameters, as well as Legendre polynomials of order 0 (constant signal), 1 (linear drift), and 2].

ROIs were functionally defined for each participant based on data from separate scanning sessions and registered to the experimental data by aligning the ROI anatomy with the experimental data using a rigid body transformation, which was then applied to the ROIs so that they were in registration with the functional data collected in the symmetry fMRI experiment.

The remainder of the analysis was performed in MATLAB. The time course data were first averaged across the three scans for each condition, and then across the voxels within each ROI. We then applied the Fourier transform to the average time course for each ROI, omitted DC, multiplied the amplitudes by 2 to get the single-sided spectrum, and scaled by dividing the amplitudes with the number of samples in the time course. An example of the resulting amplitude spectrum is shown in Figure 5. After computing the amplitude spectrum for all four conditions, we took the amplitude of two sidebands on both sides of the stimulus frequency over all four conditions, and averaging these 16 values (four sidebands \times four conditions) to get a single number representing the noise level across all conditions. We then computed the signal-to-noise ratio (SNR) for each condition by dividing the amplitude at the stimulus frequency with the noise level. We performed this procedure for each of the 12 participants to generate participant-specific SNR values for all four conditions.

Functionally defined visual regions of interest. Nine visual ROIs were defined individually for each of the 12 participants who took part in the fMRI experiments. The majority of the areas were defined based on retinotopic mapping, which was performed using the population receptive field method (Dumoulin and Wandell, 2008), with the exception of two of our participants for whom the traditional phase-mapping method with rings and wedges was used (DeYoe et al., 1994; Engel et al., 1994). The areas were defined manually based on standard criteria (Serenio et al., 1995; Engel et al., 1997): contralateral quarterfield representation for V1d, V1v, V2d, V2v, V3d, and V3v; contralateral hemifield representation for V4 (Wade et al., 2002); and V3A/B (Tootell et al., 1997). IPS0 was defined as an additional contralateral hemifield representation anterior to V3A/B (IPS0 was originally known as V7; Tootell et al., 1998; Wandell et al., 2007). VO1 as an additional contralateral hemifield representation abutting V4 (Brewer et al., 2005). We note that in three of our participants, IPS0 and/or VO1 could not be defined bilaterally, because of variability in the quality of the acquired data. Every ROI could be defined bilaterally in at least 10 participants.

In addition to the retinotopic areas, we defined LOC and motion-sensitive middle temporal cortex (hMT⁺) using fMRI block designs. During the LOC localizer scans, participants viewed blocks of images depicting common objects (12 s/block), alternating with blocks containing scrambled versions of the same objects. The stimuli were identical to the ones used in a previous study (Kourtzi and Kanwisher, 2000). hMT⁺ was defined using stimuli similar to those described by Huk and colleagues (2002): localizer scans blocks of low-contrast moving dots alternated with blocks of low-contrast static dots. Three repetitions of each of the two types of localizer scans were performed for each participant, which was sufficient to define LOC and hMT⁺ bilaterally in all participants.

EEG acquisition and preprocessing. The EEG data were collected with 128-sensor HydroCell Sensor Nets (Electrical Geodesics) and bandpass filtered from 0.3 to 50 Hz. Following each experimental session, the 3D locations of all electrodes and three major fiducials (nasion, and left and right preauricular points) were digitized using a 3Space Fastrack 3D

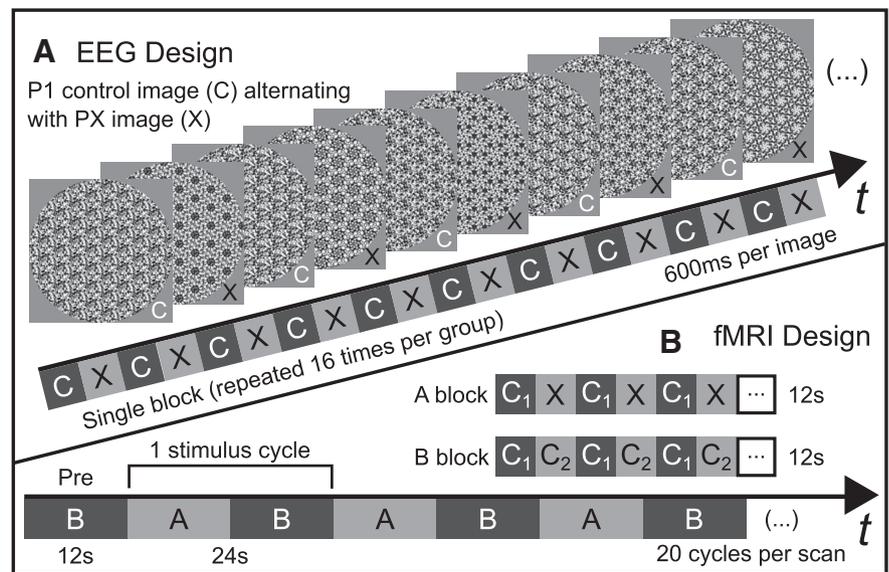


Figure 4. *A*, Schematic of the experimental design used in the EEG experiment. P1 control exemplars (marked “C”) alternated with PX exemplars (i.e., exemplars from one of the four wallpaper groups P2, P3, P4, and P6; marked “X”). Five P6 exemplars and matched controls are shown in the order they were shown in the experiment, and on a gray background like the one used in the experiment. Below the wallpaper images, a full EEG block is shown schematically, with 10 PX exemplars and 10 P1 control exemplars. *B*, Schematic of the experimental design used in the fMRI experiment. A and B blocks are shown schematically: In A blocks, P1 control exemplars (C₁) alternate with PX (X) exemplars in a way that is very similar to the EEG experiment. In B blocks, P1 control exemplars alternate with another set of P1 control exemplars (C₂). In the bottom of the figure, the first few stimulus cycles in a scan are shown schematically, including the 12 s prelude in which a B block was shown. Note that in each fMRI scan, one wallpaper group was shown exclusively, so PX was P2, P3, P4, or P6. The P1 control exemplars were directly matched to the PX exemplars (see the Control stimulus section).

digitizer (Polhemus). The digitized locations were used to construct the EEG forward model (see EEG source-imaging analysis section). Raw data were subjected to an off-line sample-by-sample thresholding procedure in which noisy sensors were replaced by the average of the six nearest spatial neighbors. On average, <5% of the electrodes were substituted; these electrodes were mainly located near the forehead or the ears, and, as such, are likely to have a negligible impact on our results, as our stimuli will likely drive responses mainly at electrodes over occipital, temporal, and parietal locations. The EEG data were then re-referenced to the common average of all the sensors and segmented into 10 1.2-s-long epochs (each corresponding to exactly one cycle of image modulation). Epochs for which a large percentage of data samples exceeded a noise threshold (depending on the participant and ranging between 25 and 50 μ V) were excluded from the analysis on a sensor-by-sensor basis. This was typically the case for epochs containing artifacts, such as blinks or eye movements.

A Fourier analysis was applied on every epoch using a discrete Fourier transform with a rectangular window. Given that each epoch was 1.2 s, this Fourier transformation led to a frequency resolution of $\delta f = 0.42$ Hz. For each frequency bin, the complex-valued Fourier coefficients were then averaged across all the epochs and all the trials. Thus, these average Fourier coefficients for each session were obtained from up to 80 data samples (10 epochs \times 8 trials). Each participant did two sessions, and the Fourier coefficients from each session were averaged immediately for the sensor space analysis. For the source localization analysis, each session was taken through the source localization pipeline separately and then averaged in source space.

EEG sensor space analysis. The average Fourier coefficients from each subject were averaged over a “sensor space region of interest” consisting of six electrodes over occipital cortex (70, 74, 75, 81, 82, 83). We used an additional second sensor space ROI, consisting of six electrodes over parietal cortex (53, 54, 61, 78, 79, 86) to generate control data. We isolated the response specific to symmetry by analyzing the odd and even harmonics of the spectrum separately (Norcia et al., 2002). Because we used a steady-state design in which P1 control images that had no rota-

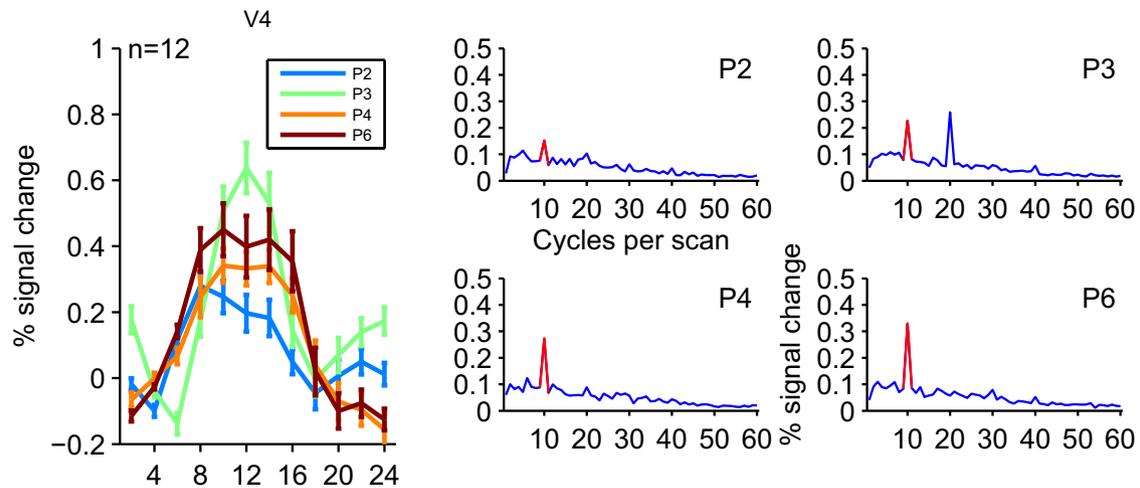


Figure 5. On the left-hand side of the figure, we plot the average fMRI BOLD response within the bilateral V4 ROI over a single stimulus cycle, averaged across participants ($n = 12$), for each of the four conditions. Error bars indicate the SEM across participants. Stimulus blocks A and B are shown in the bottom part of the plot. On the right-hand side of the figure, we plot amplitude spectra for the four conditions computed over bilateral V4 ROI, and averaged across participants. The stimulus frequency (10 cycles/scan) is indicated in red on the spectra.

tion symmetry alternated with images from the four rotation symmetry groups, the rotation symmetry response will predominantly be found in the odd harmonics, whereas the even harmonics will mostly consist of responses arising from the image update, unrelated to symmetry. To demonstrate the specificity of the odd harmonics, we collected a separate dataset where we simply alternated the same P1 control images with an additional 10 P1 images (the same additional set of control images that was used in the fMRI control block). The expectation was that there will be no signal in the odd harmonics in this case, while the even harmonics will be largely unchanged (Norcia et al., 2002).

We computed the SNR for each harmonic based on the amplitude spectrum, the noise level for each harmonic, by taking the amplitude of one sideband on either side of the harmonic frequency over all four conditions, and averaging these eight values (two sidebands \times four conditions) to get a single number representing the noise level across all conditions for that harmonic. We then computed SNR by dividing the amplitude at the harmonic frequency with the noise level. This was done individually for each participant. This procedure is analogous to the one used when computing SNR for the fMRI data, with the exception that we had to estimate the noise based on two, rather than four, side frequencies, to avoid including neighboring harmonics in the noise level.

Tissue segmentation procedure. The FreeSurfer software package (<http://surfer.nmr.mgh.harvard.edu>) was used to extract both gray/white and gray/CSF boundaries, and generate cortical surface meshes. To avoid discontinuities in the cortically constrained inversion procedure arising from curvature differences between the gray/white and gray/CSF boundary, we generated a surface partway between these two boundaries that has gyri and sulci with approximately equal curvature. This “midgray” cortical surface consisted of a very dense triangular tessellation of several hundred thousand regularly spaced vertices. This tessellation was then downsampled to 20,484 vertices using the MNE software package (<http://martinos.org/mne/stable/index.html>). This number is low enough to compute the forward model on a standard workstation and yet accurately reflect the shape of cortical manifold (see e.g., Baillet et al., 2001). The final midgray surface was used to define the visual ROIs and the source space for the EEG current modeling. The FSL toolbox (<http://www.fmrib.ox.ac.uk/fsl/>) was used to segment, from the individual T1- and T2-weighted MRI scans, contiguous volume regions for the inner skull, outer skull, and scalp. These MRI volumes were then converted into inner skull, outer skull, and scalp surfaces (Smith, 2002; Smith et al., 2004) that define the boundaries between the brain/CSF and the skull, the skull and the scalp, and the scalp and the air.

EEG source-imaging analysis. For each participant, the EEG source space was defined by the midgray surface (see the Tissue segmentation procedure section). The distance between the 20,484 vertices of this sur-

face was on average 3.7 mm (SD, 1.5 mm; range, 0.1–11 mm). Current dipoles were placed at each of these vertices. Their orientations were constrained to be orthogonal to the cortical surface to diminish the number of parameters to be estimated in the inverse procedure (Hämäläinen et al., 1993). The source space, the 3D electrode locations, and the individually defined boundaries were then combined using the MNE software package to characterize the electric field propagation using a three-compartment boundary element method (Hämäläinen and Sarvas, 1989). The resulting forward model is linear and links the activity of the 20,484 cortical sources to the voltages recorded by our EEG electrodes.

Cortical current density estimates of the neural responses were obtained from an L2 minimum-norm inverse of the forward model as described in a recent study by Cottareau et al. (2012). We used the functionally defined visual ROIs to constrain these estimates by modifying the source-covariance matrix. The aim of this procedure was to decrease the tendency of the minimum-norm procedure to smooth activity across different functional ROIs. The following two modifications were applied: (1) we increased the variance allowed within the visual ROIs by a factor of two relative to other vertices; and (2) we enforced a local correlation constraint within each ROI using the first- and second-order neighborhoods on the cortical tessellation with a weighting function equal to 0.5 for the first order and 0.25 for the second order. This correlation constraint, therefore, respects both retinotopy and boundaries between visual areas, and permits a more precise dissociation of signals from different ROIs. This is not the case for other smoothing methods, such as LORETA, that apply the same smoothing rule across all of cortex (Pascual-Marqui et al., 1994). For each participant and condition, this inversion scheme was applied to the average Fourier coefficients (see the EEG acquisition and preprocessing section), which resulted in an estimation of these coefficients for each source on the cortical tessellation. The complex coefficients were then averaged across all the sources belonging to each functional ROI, for each participant. These within-participant ROI averages were then averaged across participants to generate cross-participant averages for each ROI. This procedure has been shown to reduce cross talk between different cortical ROIs, resulting in better area resolution in the group inversion than can be achieved in single participants (Cottareau et al., 2015).

Symmetry onset-timing estimation. For the source localization analyses, our main interest was in estimating the response onset latency differences between different ROIs. We did this by fitting a piecewise linear function to the source-localized waveforms using a shape-language-modeling toolbox implemented in Matlab (<http://www.mathworks.com/matlab-central/fileexchange/24443-slm-shape-language-modeling>). The model was fit to the section of the waveform beginning at the onset of the second

image (from one of the four wallpaper groups) and ending at 300 ms later, our rationale being that the response should begin within 300 ms of the image onset. The only constraint on the model was that it had to consist of three linear sections, separated by two breakpoints that were determined by the fitting procedure. The first breakpoint provided a reasonable estimate of the latency of the negative deflection characteristic of the symmetry response (see Fig. 9).

We used a jackknife procedure to estimate the error of the latency differences between pairs of ROIs by computing average waveforms across all participants except one, and then fitting the piecewise linear model to the resulting waveform. We then subtracted the resulting response latency in the first ROI from the latency in the second. We repeated this procedure for all participants, and estimated the SE of the differences, as described by Miller et al. (1998, their Eq. 2).

Computational modeling. Wallpaper patterns can be construed as being textures or, alternatively, as having elements of form that relate to object properties. There has been considerable interest in developing computational models of early visual cortical mechanisms or machine vision algorithms that could support either texture or object processing. When interpreting our findings, it is important to assess the extent to which differences in the neural responses elicited by the four wallpaper groups can be explained by properties of the stimuli that are unrelated to rotation symmetry, per se, but could be related to other types of regularity found in textures. To address this question, we measured the predicted strength of neural responses to the four wallpaper groups according to (1) their Weibull statistics, (2) their discriminability by a robust texture classifier, (3) their pattern phase congruency across spatial scale, and the (4) the responses of the model of second-order contrast sensitivity. Each of these approaches has been shown to be sensitive to the statistics of natural images and can provide quantitative metrics of texture similarity. Our computational approach was analogous to the one taken in the EEG and fMRI experiments, in that we assessed the predicted response difference between each group of exemplars and its corresponding P1 control. All four models were run on the same set of images that was used for the EEG and fMRI experiments, with 20 exemplars per wallpaper group, and the predicted responses differences were then averaged across exemplars within a group. The first three approaches were applied to cropped versions of the stimuli to reduce the number of artifacts caused by the aperture, while the contrast model was applied to the whole image, including the aperture, and the resulting output map was then cropped.

It has been suggested that parameters of the Weibull distribution provide useful summary statistics for natural images (Geusebroek and Smeulders, 2002). The Weibull distribution was initially used to describe the range of particle sizes resulting from roller mills (Rosin and Rammler, 1933), and it has intuitive appeal as an index of “fragmentation” in an image (Geusebroek and Smeulders, 2002; Groen et al., 2013). The Weibull distribution is parameterized by a “scale” parameter a and a “shape” variable b . The scale parameter varies with the contrast energy of the image, and the shape parameter varies with the spatial coherence of the image (Groen et al., 2013). Several recent event-related potential (ERP) reports have investigated the extent to which Weibull statistics could explain the variance in the neural responses (Scholte et al., 2009; Groen et al., 2012a,b; Groen et al., 2013). The Weibull statistics capture a substantial fraction of the ERP variance in response to natural images that vary widely in their structure. We computed the Weibull statistics (a and b) for each image, $I(x,y)$, as follows: First, the image was convolved with first-order Gaussian derivative filters for some Gaussian $G(x, y; \sigma)$. We then computed the gradient magnitude over the image as follows:

$$\nabla^1 I(x, y; \sigma) = \sqrt{\left(I(x, y) \otimes \frac{\delta G(x, y; \sigma)}{\delta x} \right)^2 + \left(I(x, y) \otimes \frac{\delta G(x, y; \sigma)}{\delta y} \right)^2}$$

Then, a 256-bin histogram was calculated over $\nabla^1 I$, which was modeled using a Weibull distribution, as follows:

$$f(x, a, b) = \frac{b}{a} \left(\frac{x}{a} \right)^{b-1} e^{-\left(\frac{x}{a} \right)^b},$$

using the wblfit command in MATLAB. To maximize the fit between the Weibull distribution and our data, a and b were calculated for a range of spatial scales ($2 \leq \sigma \leq 6$), and we used the scale that gave the highest (logistic regression) accuracy when classifying the exemplars from their matched P1 controls. Finally, we used β_a and β_b to scale a and b and construct a distance metric between pairs of images I_1 and I_2 , as follows:

$$d(I_1, I_2) = \sqrt{\beta_a(a_1 - a_2)^2 + \beta_b(b_1 - b_2)^2},$$

where a_i, b_i are the Weibull statistics from image I_i . This distance metric was generated for each exemplar and its matched P1 control.

Our second approach was to ask whether features identified as informative by the computer vision and texture classification literature can predict the neural responses to our symmetry patterns. A recent study (Dong et al., 2014) evaluated 51 computational texture features in terms of how well they could rank patterns in terms of their visual similarity and found that classification based on local image-patch features (Varma and Zisserman, 2003) offered the best performance. The local image-patch approach is also attractive because it shares several similarities with the sparse coding model of Olshausen and Field (1997). In particular, both models are based on extracting $n \times n$ pixel local neighborhoods (“textons”) from a set of images and then using an unsupervised clustering algorithm to learn a set of textons that can be used to represent the images. The main difference between approaches is the algorithm used to learn the texton dictionary, although both methods learn dictionaries of a similar size. Varma and Zisserman (2003) used dictionaries consisting of 1000–3000 exemplar textons, while Olshausen (2013) has recently experimented with 10-fold overcomplete sparse coding dictionaries consisting of 2048 basis functions. Critically, as noted above, features from the local patch model were among the best predictors of human perceptual similarity ratings for a large texture database (Clarke et al., 2011a). Our implementation involves downsampling the image by a factor of four, and then using k -means to learn a dictionary of ($k = 100$) $n \times n$ local neighborhoods ($n = 3$). Once a texton dictionary is created, an image, I_p , can be represented by a histogram, h_p , over the textons, and two textures are compared with one another using a simplified form of the Bhattacharyya distance, as follows:

$$d(h_1, h_2) = 1 - \sqrt{h_1 h_2}.$$

Using this approach, we calculated the distance between each exemplar and its corresponding P1 control.

Natural images contain numerous high-order correlations, among which correlations in the phase spectrum across scale are particularly prominent. Several models of human feature perception have been built to extract features in the frequency domain that are due to the presence of various types of edges that can create phase correlations. Inherent in these models is the extraction of correlations in spatial phase across scale (Morrone and Burr, 1988; Owens, 1994; Kovess, 2000). Perceptual appearance is strongly controlled by cross-scale phase coherence (Morrone and Burr, 1988), and sensitivity to certain changes in the appearance of a texture can be accurately modeled ($R^2 > 0.95$) using the variance of the phase congruency map (Emrith et al., 2010). We tested whether patterns of phase coherency across scale could predict the differences in neural responses between the wallpaper groups, by computing a summary difference on phase coherency in terms of the log of the variance of the phase congruency map between each exemplar image and its matched P1 control.

The second-order contrast (SOC) model is a cascaded, feedforward model of BOLD responses to contrast and texture, in which the image is first processed through a bank of contrast-normalized, localized, V1-like filters, the output of which is then reprocessed by a second stage that, among other things, measures the contrast variability in the output of the first stage (Kay et al., 2013). The SOC model has been shown to be effective at predicting differences in texture responses between V1 and V2/3 (Kay et al., 2013), so it is a plausible candidate model for predicting differences in the response to the different wallpaper groups. Furthermore, the model is attractive as a test case because second-order contrast is higher in natural images with intact phase spectra (Kay et al., 2013), and because the model has been fit to data for V1, V2, V3, and V4,

successfully capturing increased sensitivity to the structure of natural scenes in the higher-order areas. We applied the SOC model separately to each exemplar image, using model parameters for V1, V2, V3, and V4 provided in the original publication (Kay et al., 2013). We then summed the model output across each image, and computed the modeled difference between each exemplar and its matched P1 control.

Results

Parametric modulation of fMRI BOLD signal with order of rotation symmetry

We saw a strong and consistent response to rotation symmetry in four of our functionally defined ROIs: V3, V4, VO1, and LOC. Plotting the average fMRI BOLD signal over a single cycle of the stimulus for the four wallpaper groups gives us an idea of what the response looked like (Fig. 5, left). Importantly, this single-cycle average reveals a response that begins 4–6 s after the onset of the A block (in which P1 control exemplars alternated with PX exemplars; Fig. 4) and returns to a lower level around 4–6 s after the onset of the B block (in which only P1 exemplars are shown). This response profile is consistent with the PX exemplars eliciting a symmetry response above and beyond that being elicited by the P1 exemplars. Because of the hemodynamic delay (Boynton et al., 1996), it is expected that this response would begin several seconds after the onset of the A block.

The spectral analysis confirms this interpretation. There is a clear peak at the stimulus frequency in the amplitude spectra for all four groups (Fig. 5, right, V4). There is another peak at the second harmonic of the stimulus frequency, but only for wallpaper group P3. The four wallpaper groups are extremely well matched in our experimental design, so it is unclear why only P3 would produce this second harmonic response. The second harmonic is indicative of transient responses occurring at the transitions between the A and B blocks. This can be seen in the temporal domain, where the second harmonic shows up as a second peak at $\sim 0/24$ s. Because only P1 stimuli were shown during the B block, the second harmonic is not specific to symmetry. By contrast, the first harmonic captures the sustained response to the presence of symmetry during the A block and the lack of it during the B block. We therefore focus on the first harmonic for the remainder of the analysis.

Several ROIs had a very similar response pattern to V4 (V3, VO1, LOC), while others had very little response to symmetry (MT, IPS0), as indicated by flat single-cycle averages around 0% signal change and spectral peaks at the stimulus frequency that are not different from the neighboring frequencies (Fig. 6). We can better visualize the profile of sensitivity to the different symmetry groups across ROIs by computing the SNR at the stimulus frequency (see Materials and Methods for details). The four ROIs that show a strong response to symmetry all have a very similar profile: SNR increased parametrically with the maximum order of rotation symmetry contained within the group (see Fig. 7). We evaluated these effects with a repeated-measures ANOVA (rANOVA) and a linear mixed-effects analysis (LMEA) implemented by the lme4 (Baayen et al., 2008) and lmerTest (Kuznetsova et al., 2014) packages in R (R Core Team, 2014). To test for overall effects, we first performed a rANOVA with ROI as the first factor, condition as the second factor, hemisphere as the third factor, and hemisphere-specific SNR values as the dependent measure. We found a main effect of ROI ($F_{(8,757.3)} = 49.1, p < 0.0001$), a main effect of condition ($F_{(3,757.0)} = 51.6, p < 0.0001$), and a significant interaction between ROI and condition ($F_{(24,757.0)} = 4.0, p < 0.0001$). There was no significant main effect of hemisphere ($F_{(1,757.0)} = 1.6, p =$

0.21), and all other interactions were far from significance ($p > 0.25$). This result indicates that there is a strong effect of rotation symmetry that varies among ROIs, but not systematically between hemispheres.

To specifically test for the presence of a linear effect of maximum order of rotation symmetry contained in the wallpaper groups within each of our nine ROIs, we then performed an LMEA that tested for a linear relationship between SNR and condition. We disregarded hemisphere and used SNR values computed across all voxels within each bilateral ROI, as the dependent variables. We defined ROIs and conditions as fixed effects, and participants as a random effect. We applied contrasts looking for linear effects of condition within each ROI, coding MT as the base level relative to which our the effects in other ROIs were judged, under the assumption that there is minimal symmetry coding in MT. The linear contrast was significant in areas V3 ($t_{(387)} = 2.49, p = 0.013$), V4 ($t_{(387)} = 4.46, p < 0.0001$), VO1 ($t_{(387)} = 4.57, p < 0.0001$), and LOC ($t_{(387)} = 4.13, p < 0.0001$). For all other tested ROIs, the contrast was far from significance ($p > 0.25$). The p values and denominator degrees of freedom for both analyses were calculated using Satterthwaite's approximations (Kuznetsova et al., 2014). This result indicates that in a subset of the ROIs (V3, V4, VO1, and LOC) the response to symmetry increased parametrically with the maximum order of rotation symmetry within the patterns.

Parametric dependence of visual-evoked potentials on order of rotation symmetry

We first demonstrate that the odd harmonic components of the EEG visual-evoked potential reflect neural responses specific to the rotation symmetry in the images. The top half of Figure 8 shows the raw cycle averages for responses to alternations of P1 control images (A) and P1 control and P6 test images (B), within our occipital sensor space ROI (described in more detail in the EEG sensor space analysis section of Materials and Methods and indicated with black rings on the second topographic plot in Fig. 9). The two waveforms are clearly different upon visual inspection. To gain analytic insight into the components that are specific to symmetry, we generated synthesized waveforms by isolating the odd and even harmonics in the spectral domain, and back-transforming them separately into the time domain. The resulting waveforms consist exclusively of the signal from the odd harmonics and even harmonics, respectively (plotted in Fig. 8C,D). We can now compare the waveforms from the control experiment with those from the main experiment. For the sake of simplicity, we consider only wallpaper group P6 and the P1 controls generated from P6 images. In both experiments, there is a large positive peak in the even harmonic waveforms beginning ~ 100 ms after the onset of both the first and second presented image, which is followed by a large negative peak at ~ 300 ms.

In the control experiments, there is no measurable signal in the odd harmonic waveforms, but in the main experiment there is a sustained, nearly square-wave, negative-going response beginning ~ 100 ms after the onset of the second image (which always came from rotation symmetry group P6). This negative response is sustained until ~ 100 ms after the transition from P6 back to P1. The fact that there is no signal in the odd harmonic when P1 images alternated with other P1 images demonstrates that the signal in the odd harmonics is exclusively related to symmetry. This is not true in the even harmonics, where there is clearly a signal in both experiments, presumably driven by local contrast transients. We do see a small difference between the even harmonic waveforms from the two experiments; previous mod-

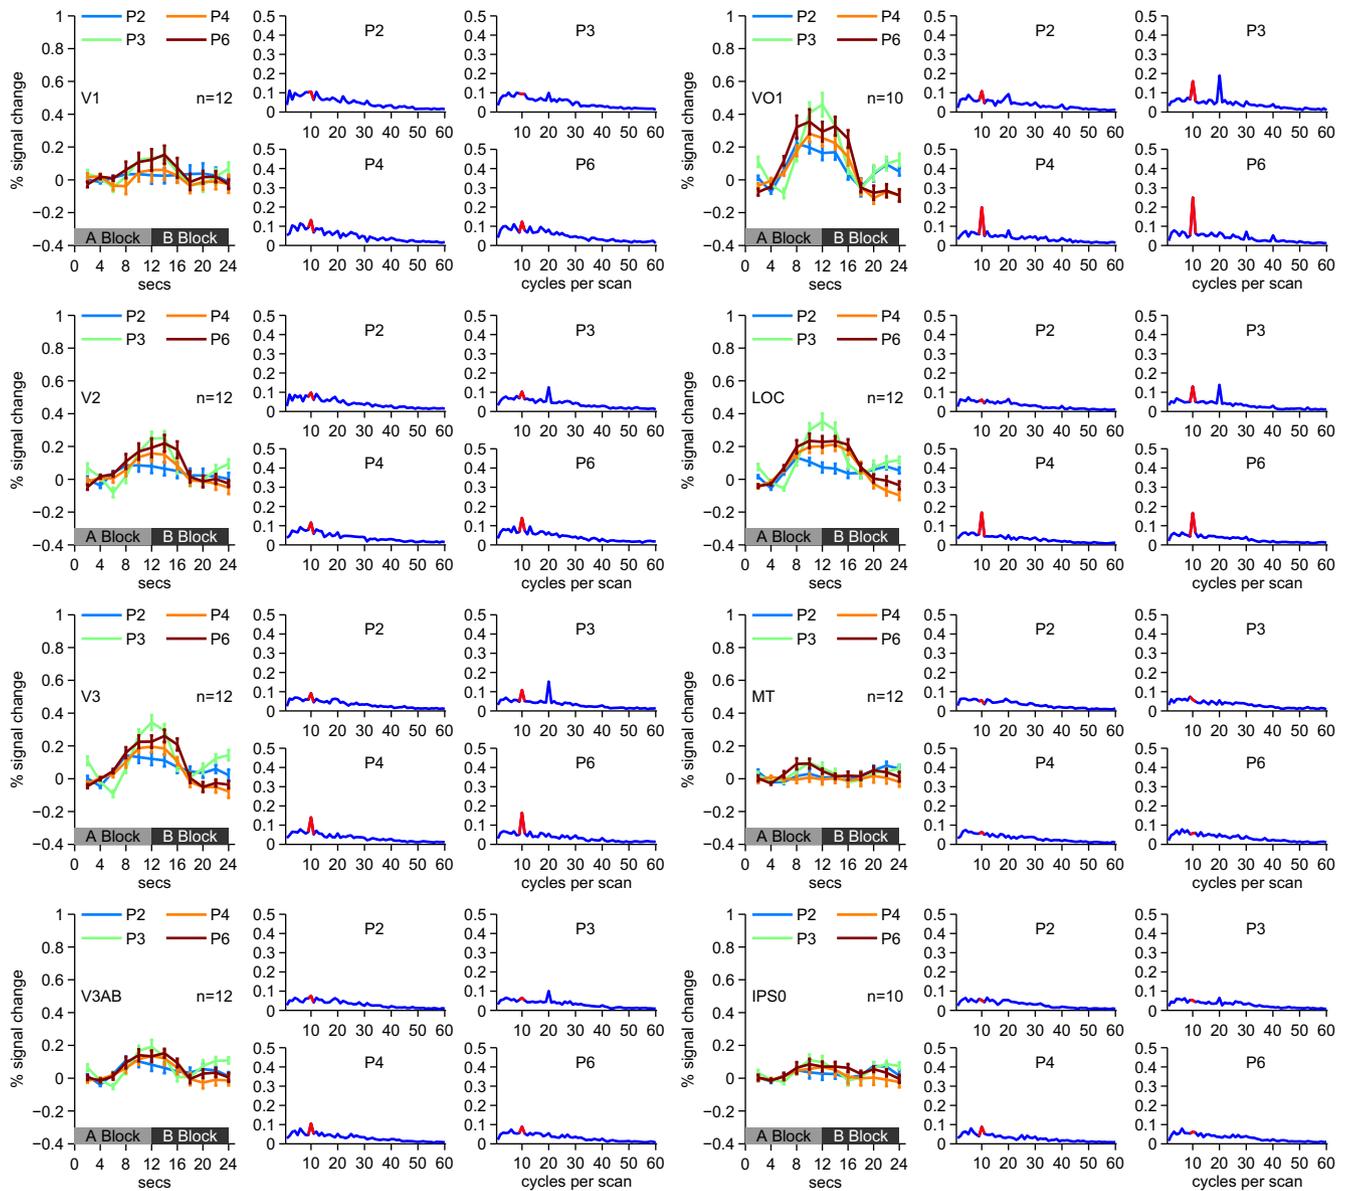


Figure 6. Average fMRI BOLD responses over a single stimulus cycle and average amplitude spectra for the four conditions computed over eight bilateral ROIs (excluding V4). The subplots in this figure follow the same conventions as in Figure 5, which show the same data for the bilateral V4 ROI.

eling work has demonstrated that when two stimuli are alternating, and eliciting different neural responses, an increase in the response to just one of the stimuli can lead to greater response amplitudes in both the even and odd harmonics (Cottareau et al., 2011). Nonetheless, our comparison of the main experiment and the P1 control experiment clearly shows that the odd harmonics are entirely specific to symmetry, and that the even harmonics are mostly nonspecific. Because of this, all further EEG analyses in this report focus exclusively on the odd harmonics.

We limited our analysis to the first four odd harmonics, under the assumption that most of the signal would be found in the lower harmonics. We computed the SNR for each of the first four odd harmonics, separately for each of the four conditions, within the occipital sensor space ROI (Fig. 9, black rings on the second topographic plot). The results nicely replicated our fMRI analysis, in that the symmetry response measured by sensors over occipital cortex increased parametrically with the maximum order of rotation symmetry within each pattern. This sensitivity to ro-

tation symmetry was specific to the first and third harmonics of the stimulus frequency.

Our statistical analysis of the EEG sensor space data took an analogous approach to the one used for the fMRI data, using the same R packages. We first performed an rANOVA with harmonic as the first factor, condition as the second factor, and the EEG SNR as the dependent measure. We found a main effect of harmonic ($F_{(3,360.0)} = 6.98, p = 0.00014$), a main effect of condition ($F_{(3,360.0)} = 20.0, p < 0.0001$), and no significant interaction ($F_{(9,360.0)} = 1.6, p = 0.10$); and all other interactions were far from significance ($p > 0.25$). To explore the response to rotation symmetry for each harmonic, we then did an LMEA, applying contrasts testing for linear effects of condition for each harmonic. As the base level relative to which the effects were judged, we used the SNR computed for the first harmonic measured in the parietal control ROI (described in Materials and Methods; Fig. 9, white circles on the third harmonic scalp map). This was done under the assumption that there would be little or no sensitivity

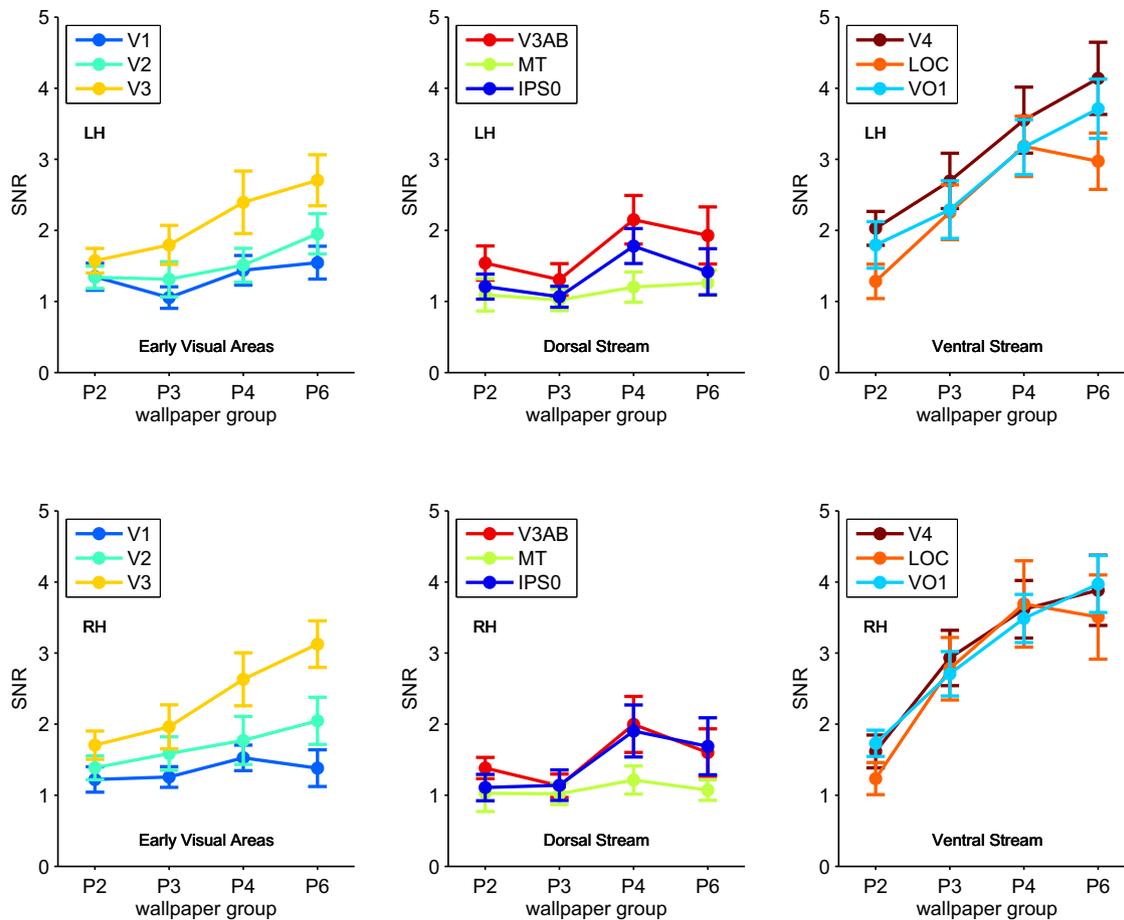


Figure 7. Responses to the four wallpaper groups across nine different functionally defined ROIs. The y-axis indicates the SNR computed at the stimulus frequency, and the x-axis indicates the four wallpaper groups. Left hemisphere ROIs are plotted in the top part of the figure, whereas right hemisphere ROIs are plotted in the bottom part. For visualization purposes, we have split the ROIs into the following three groups: early visual areas (left side), dorsal stream areas (middle), and ventral stream areas (right side). Error bars indicate the SEM.

to rotation symmetry in parietal cortex. The linear contrast was significant for the first ($t_{(465,9)} = 4.10, p < 0.0001$) and third harmonic ($t_{(465,9)} = 3.34, p = 0.0009$), but not for the fifth ($t_{(465,9)} = 1.57, p = 0.12$) or seventh harmonic ($t_{(465,9)} = 1.04, p = 0.30$). To make sure that our results were not an artifact of our chosen base level, we also repeated the LMEA using the SNR computed for the third, fifth, and seventh harmonics, respectively, from the same control ROI as an alternative base level, but the results were consistent across all four different base levels.

Temporal order of symmetry processing in functionally defined ROIs

Our source space analysis of the EEG data focused on a subset of the ROIs, all of which had a parametric response in the fMRI analysis: V3d, V4, and LOC. Because ventral V3 is immediately adjacent to V4 and therefore likely to be poorly resolved from V4, we limited our analysis to the dorsal part of V3. Similarly, VO1 was left out of the analysis due to its likely confusability with V4. The source-localized EEG results were consistent with the fMRI results: The V3d, V4, and LOC ROIs each had odd-harmonic responses to symmetry (Fig. 10). The main power of a source-localized EEG, however, is that it gives us insight into the temporal order of processing in the different ROIs, so that we can track the development of the symmetry response across several areas of the visual cortex. We found that the processing of symmetry begins in V3 and V4 as early as 75 ms after a symmetry image is presented, and, after a delay of ~30 ms, the LOC response begins.

The short response latency makes it unlikely that the responses in V3d and V4 are the result of feedback from areas later in the visual processing pathway, and instead suggests that symmetry perception relies on a bottom-up process that begins in V3d and V4, and that the responses in LOC rely on input from V3d and V4.

We determined the temporal order by estimating the onset latency of the symmetry response in each ROI, based on synthesized waveforms for each of the four groups, which consisted exclusively of signal from the first six odd harmonics. We focused on the first six odd harmonics, because our sensor space analysis had indicated that most of the symmetry response was found in the lower harmonics, so the higher harmonics would most likely only add noise to our onset estimates. We averaged these waveforms across the four groups, for each participant, because we were interested in characterizing the general response to rotation symmetry within each ROI. Based on these waveforms, we then estimated the onset latency differences between every pairing of the three ROIs, and estimated the SE of the latency differences among the ROIs using a jackknife procedure (for details, see Materials and Methods).

For illustration purposes, we also used the same procedure to compute the response latency for each individual ROI, based on average waveforms across all participants, which yielded the following latencies: V3d, 75 ms; V4, 77 ms; and LOC, 109 ms. We then computed the SE of these latencies using the same jackknife procedure as was used for the latency differences, which was used to generate error bars (Fig. 10). We computed *t* statistics for a two-tailed

test comparing each ROI pair by dividing the onset difference estimate, based on the overall sample with the SE of the difference produced by the jackknife method (Miller et al., 1998). The onset latencies were not significantly different between V3d and V4 ($t_{(12)} = 0.22, p = 0.83$), but LOC had a significantly later onset than both V3d ($t_{(12)} = 2.67, p = 0.02$) and V4 ($t_{(12)} = 3.21, p = 0.007$).

Computational models that are unable to predict empirical results

We applied four different computational approaches to determine whether the operations that are instantiated by them are sufficient to code for the presence of rotation symmetry, and, more specifically, whether their output matches the parametric pattern of neural responses to the four wallpaper groups seen in the EEG and fMRI experiments. For each modeling approach, we measured the predicted response differences between the exemplars from wallpaper groups P2, P3, P4, and P6, and their matched P1 controls. All four approaches were able to distinguish the exemplars from the controls, but none predicted the parametric response to rotation symmetry that was indicated by our empirical data.

We computed Weibull statistics over a range of spatial scales (from 2 to 5 pixels) and defined the optimal value of σ as the value that maximized the accuracy of a logistic regression classifying the exemplar images from their corresponding control. In our case, the optimal σ value was 2.5 pixels, which resulted in coefficients of $\beta_a = -517$ and $\beta_b = 0.116$, and mean classification accuracy across the four wallpaper groups of 0.69. The difference metrics that were generated based on the coefficients of the logistic regression for each exemplar and its matched P1 control are shown in Figure 11A, averaged across exemplars for each wallpaper group. These metrics do not differ among the four symmetry groups, indicating that Weibull statistics do not vary systematically between different symmetry groups. Weibull statistics are thus unlikely to underlie the parametric responses to rotation symmetry revealed by our EEG and fMRI experiments. They could, however, contribute to differences between responses evoked by the exemplars from the four wallpaper groups (P2, P3, P4, and P6), and their matched P1 controls, and as such indicate the presence or absence of some degree of rotation symmetry.

Our texture classification approach returned a histogram distance between each exemplar and its corresponding P1 control, averaged across exemplars within each wallpaper group, which is plotted in Figure 11B. The fact that the values are above zero for all four wallpaper groups indicates that this measure is able to discriminate exemplar images with rotation symmetries (P2, P3, P4, and P6) from their matched P1 controls, and, like the Weibull statistics, this measure could distinguish the presence or absence

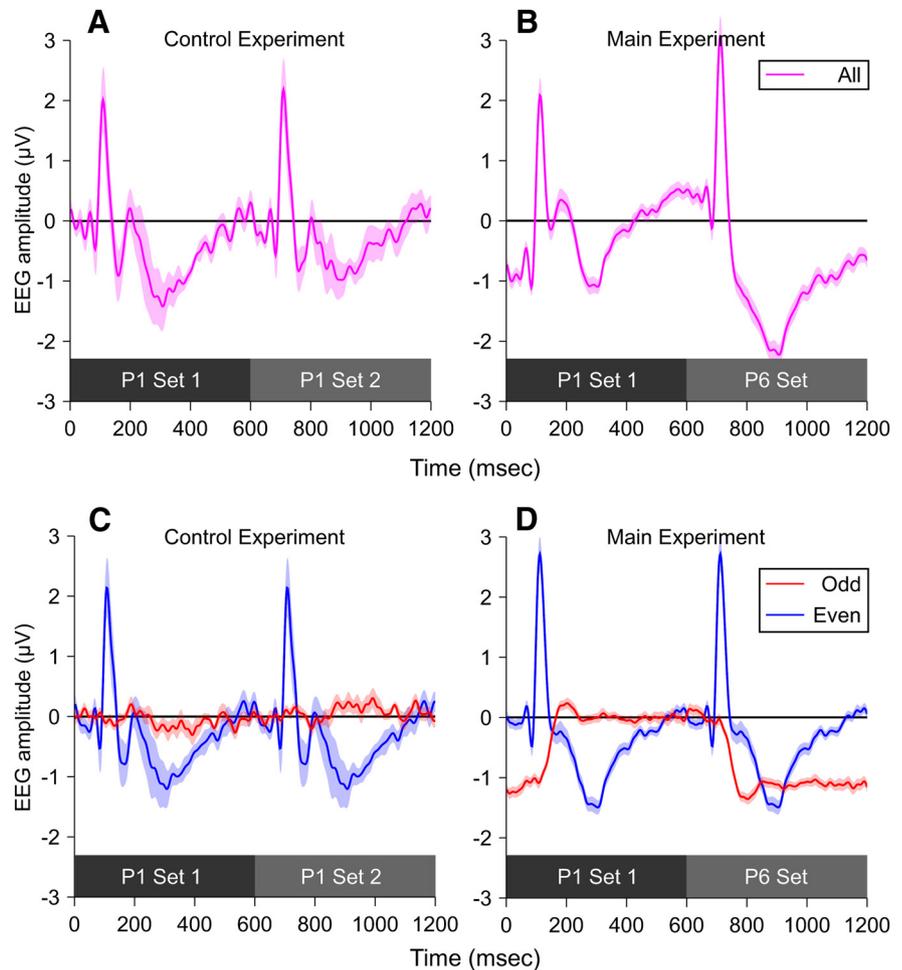


Figure 8. A–D, EEG waveforms from a control experiment in which P1 exemplars alternated with another set of P1 exemplars (A, C), compared with waveforms from our main EEG experiment, in which P1 exemplars alternated with P6 exemplars (B, D). Plots A and B show unfiltered waveforms, while C and D are waveforms that have been split into the odd and even harmonics in the Fourier domain, and then back-transformed into the time domain. The even harmonic waveforms (blue traces) are quite similar between the two experiments, whereas the odd harmonics (red traces) look quite different between the experiments, indicating that the additional neural processing evoked by the P6 exemplars is captured mostly by the odd harmonics. All waveforms were generated separately for each participant, and averaged across participants; the shaded areas indicate the SEM.

of rotation symmetry. The histogram distances are larger for P3 and P6, compared with P2 and P4, indicating that the local patch statistics are unable to capture the parametric dependence of the responses to rotation symmetry found in the fMRI and EEG experiments.

The phase congruency difference between each exemplar and its matched P1 control, averaged across exemplars within each wallpaper group, is plotted in Figure 11C. For some of the groups, P3 especially, this measure picks up on the difference between the exemplars and the control stimuli, indicating sensitivity to the presence or absence of rotation symmetry, but the magnitude of the phase congruency difference does not vary parametrically in the way that we would predict given the fMRI and EEG data.

The average differences in the summed outputs of the SOC model between each exemplar and its matched control are plotted in Figure 11D. The fact that all of the responses are positive indicates that the model always predicted a larger response to the exemplar compared with its matched control, which again indicates some sensitivity to the presence or absence of rotation symmetry. However, the model output resembles that of the local image-patch metric (Fig. 11, compare B, D), regardless of which

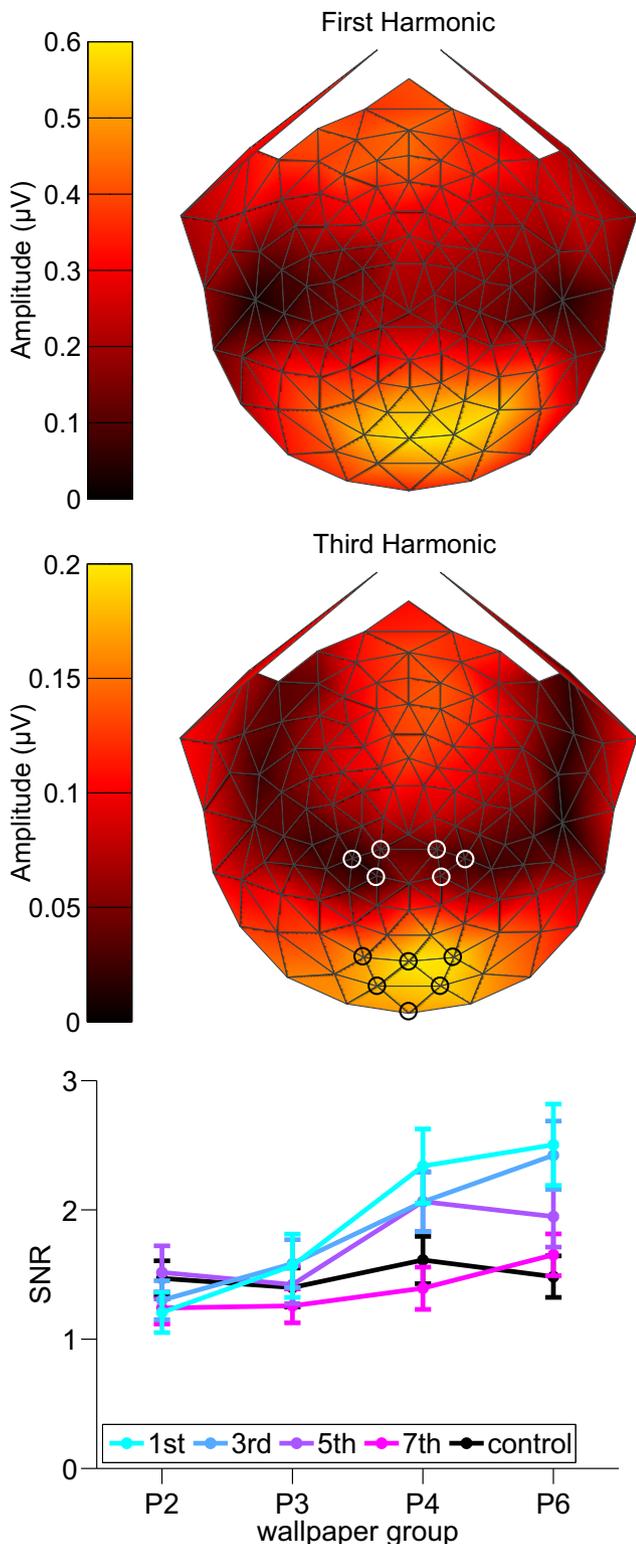


Figure 9. Sensor space results. The topographic plots show the signal in the first and third harmonics of the stimulus frequency, averaged across the four wallpaper groups and across participants ($n = 25$). The black rings on the topographic plot for the third harmonic indicate the ROI, for which the main result was computed, consisting of sensors over occipital cortex. The white rings indicate a control ROI, consisting of sensors over parietal cortex. The bar plot in the bottom third of the figure show SNR for each of the first four odd harmonics, in the occipital ROI. The SNR has been averaged over participants for each of the four groups. Error bars indicate the SEM. The black line on the plot indicates the SNR for the first harmonic in the parietal control ROI.

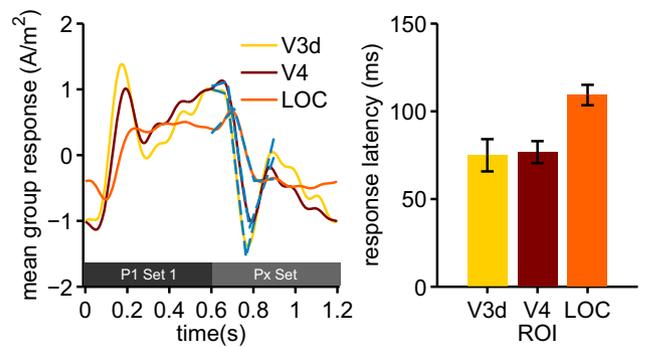


Figure 10. The timing of the symmetry response in three source-localized regions of interest. The left side of the figure shows time courses of the EEG signal within three source-localized ROIs, averaged across the four wallpaper groups and across participants. In the first 600 ms, the P1 exemplar images were shown, and in the last 600 ms the matched exemplars from the four wallpaper groups were shown. The blue lines indicate the piecewise linear models fitted to the average time courses. The right-hand side of the figure shows the estimated delay in the onset of the symmetry response, in units of milliseconds following the onset of the second image. The plotted response onsets were estimated based on the mean response across all participants, and error bars show the SEs of these estimates, generated using a jackknife method.

parameters were used, and thus does not capture the parametric responses to rotation symmetry that we saw in the fMRI and EEG data. The model output also fails to capture the within-area differences seen in the fMRI data: the predicted V4 sensitivity to symmetry in area V4 is not larger than the sensitivity in area V1, despite the obvious differences between the two areas seen in Figure 7.

Discussion

The results of our fMRI experiment demonstrate that rotation symmetry in wallpaper groups is parametrically represented in a network consisting of visual areas V3, V4, VO1, and LOC. The source-localized EEG analysis indicates that the symmetry-specific responses in V3 and V4 precede responses in LOC, suggesting that wallpaper responses in V3 and V4 are not due to feedback from later areas.

A critical aspect of our results is that the parametric response to rotation symmetry is seen in V3, but not in V1 and V2. This means that the wallpaper stimuli allow us to make a clear functional distinction between V3/V4 and V1/V2, something that has been notoriously difficult to do. Since the foundational work of Hubel and Wiesel (1959) in characterizing the response properties of V1 and V2 neurons, a theoretical framework has developed in which each successive processing stage beyond V1 is proposed to be sensitive to more complex stimulus properties. Indeed, some stimulus characteristics appear to be represented by V2, but not V1, such as relative rather than absolute retinal disparity (Thomas et al., 2002), stereoscopic edges (von der Heydt et al., 2000), 3D surface configurations (Bakin et al., 2000), and border ownership (Zhou et al., 2000). Many of these properties are also represented beyond V2, however, and studies of single-unit shape and curvature representations in V2 and V4, perhaps more closely related to the present work, have found that it is rare to identify response properties that are encountered in V4, but are absent or sparse in V2 (Hegd  and Van Essen, 2007). This pattern was repeated in fMRI data from Freeman et al. (2013), who demonstrated that V2, V3, and, to a lesser extent, V4 were sensitive to the natural texture stimuli they used, while V1 was not. Once again, V2 could be functionally separated from V1, but V2 could not be distinguished from later areas. These results underscore the significance of our ability to functionally distinguish V3 and

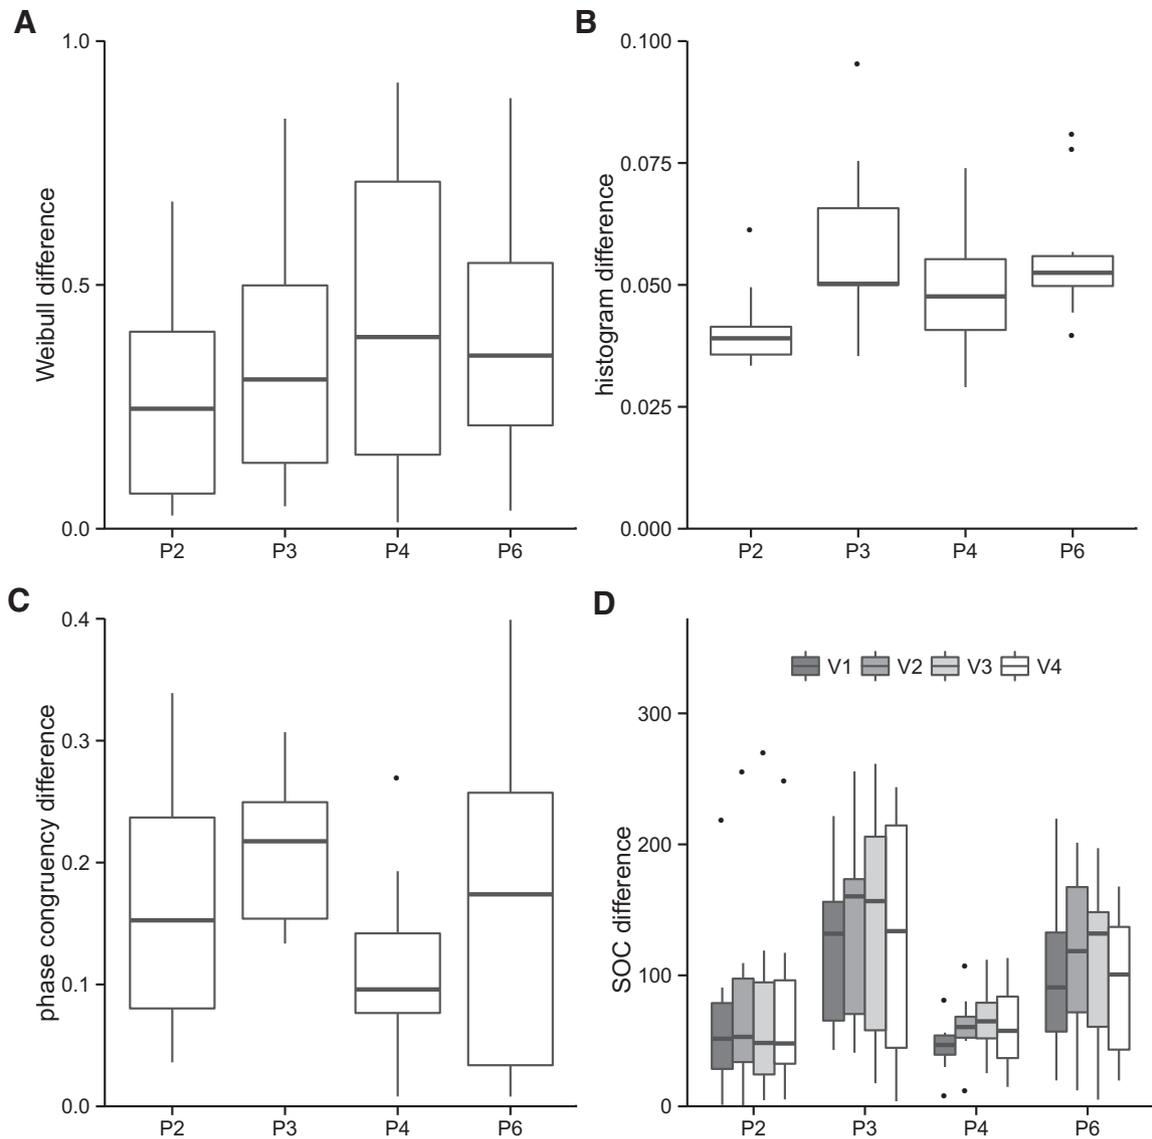


Figure 11. The results of the four computational modeling approaches we applied to the wallpaper images used in our experiments, plotted as standard boxplots. **A**, Difference metrics between each exemplar and its matched P1 control based on the coefficients of the logistic regression based on the Weibull statistics. **B**, Histogram difference between each exemplar and its corresponding P1 control averaged across exemplars within each wallpaper group. **C**, Phase congruency difference between each exemplar and its matched P1 control. **D**, Average differences of the summed outputs from the SOC model between each exemplar and its matched control. All computational modeling approaches fail to reproduce the parametric response to rotation symmetry that is seen in the neural data.

later areas from V2 and V1. The presence of both rotation symmetry and periodicity (translation symmetry) makes the wallpapers maximally regular, differentiating them from conventional textures (Portilla and Simoncelli, 2000). Our results indicate that at least under the conditions of our measurements, maximal regularity constitutes an additional level of complexity that can be encoded by V3 and beyond, but not earlier areas.

We see a clear distinction between dorsal and ventral stream areas in our results. Each ventral stream area we studied showed a strong parametric response to symmetry, and, in contrast, none of the three areas that can be considered part of the dorsal stream showed such a response. In area MT, specifically, we saw no measurable response to any of the four groups, while IPS0 and V3AB both had a response to P4 and perhaps to P6, but little or no response to the other groups. In the ventral stream, VO1 and V4 had the strongest parametric response of all our ROIs, and very similar responses to each other. LOC was very similar to V4 and

VO1 except that the response to P6 was weaker. The weaker dorsal stream responses to the static and slowly (0.8 Hz) updating wallpaper stimuli could be due to differences between the two streams that are unrelated to symmetry, such as temporal frequency tuning and preference for moving stimuli. Further work will be needed to determine whether dorsal areas are more responsive to wallpaper symmetry under different stimulation conditions.

The pattern of results in the present study is broadly similar to that found in previous fMRI work on reflection symmetry by Sasaki et al. (2005), but there are important differences. We find clear evidence of parametric responses in V3 that were very similar to those in later ventral stream areas, while Sasaki et al. (2005) found a V3 response to reflection symmetry only during passive viewing, and not when using an attention control task similar to ours. Moreover, V3 and MT had nearly identical responses in their data, while we found no response in MT.

We also find VO1 to have a strong response to rotation symmetry. VO1 has been little studied, and much of the previous work on VO1 and its sister area VO2 has focused on their role in color processing (Brewer et al., 2005; Brouwer and Heeger, 2009), although there is also evidence that VO1/2 encodes information about objects (Arcaro et al., 2009; Vandembroucke et al., 2014). Our results demonstrate that VO1 has a specific and detailed representation of texture that is entirely unrelated to color processing and is clearly differentiated from previous object-processing results.

V3 is the earliest area in the visual processing stream that has a parametric response to rotation symmetry. V3 is usually considered to be part of the early visual cortex, occupying a processing stage before the split between ventral and dorsal streams. We find it interesting that the sophisticated form processing necessary for a parametric representation of the wallpaper groups is present at such an early stage of visual processing, especially since our source localization results indicate that it is unlikely that responses in V3 are due to feedback from later areas. The response properties of V3 have not been well characterized, and it is often not included in models of the ventral stream (Riesenhuber and Poggio, 2000; Bar et al., 2001; Deco and Rolls, 2004; DiCarlo and Cox, 2007), perhaps because macaque V3 has a much smaller surface area than nearby areas like V1 and V2 (Burkhalter et al., 1986; Gattass et al., 1988). In humans, however, V3 is larger (Tootell et al., 1997; Dougherty et al., 2003), and, because of this, it has been suggested that it may play an important role in visual processing (Wandell and Winawer, 2011). We now show that this is indeed the case.

A limitation of the fMRI results presented here, as well as previous fMRI results on symmetry (Sasaki et al., 2005), is that they cannot speak to the temporal order in which the areas are activated. This means that activations in areas earlier in the hierarchy could arise either from feedback or feedforward connectivity, or from a combination of the two. We addressed this issue with source-localized EEG and showed that the response to symmetry in V3 and V4 leads the response in LOC by ~ 30 ms. This finding, along with the short-onset latencies in V3 and V4 (~ 75 ms), suggests that representations of symmetry in these areas arise through a feedforward process, rather than being due to feedback from higher areas such as the LOC. To definitively exclude contributions from feedback activity, however, it will be necessary to measure response timing under open-loop conditions (i.e., with an evoked transient event-related design with random intertrial intervals), which is preferable to the closed-loop conditions that were imposed by the steady-state design used in the current EEG experiment.

The role of V3 and V4 in representing rotation symmetry is particularly interesting given recent transcranial magnetic stimulation work suggesting that LOC plays a causal role in the processing of reflection symmetry (Bona et al., 2014). Our data do not contradict this result, but do suggest that, at least for rotation symmetry, other areas may contribute to the symmetry processing in LOC. A potential avenue for resolving this question is the combination of source-localized EEG and response-locked analysis techniques, which have successfully been applied to investigating the causal role of different visual areas in other perceptual domains (Ales et al., 2013; Cottareau et al., 2014).

The evidence presented here constrains computational models of symmetry processing—the models must be able to capture the parametric response to rotation symmetry. We tested the ability of four well developed feedforward models to account for our results by presenting them with the images used in our ex-

periments. While some did a good job of detecting rotation symmetry (i.e., distinguishing test and control images), none produced the parametric responses we see in the empirical neural data. An important challenge for future work will be to generate new models that incorporate the known properties of neurons in early visual areas and that can predict the parametric response patterns. Modeling can also be used to determine whether or not feedback is required for such specificity or instead can be derived in the purely feedforward fashion, as suggested by our EEG results.

Because receptive field size varies greatly among visual cortical areas and tends to increase as one ascends the visual processing hierarchy (Harvey and Dumoulin, 2011), it is important to consider whether receptive field size might play a role in our observed symmetry sensitivity. If the differences in sensitivity to symmetry we see between areas were driven exclusively by differences in receptive field size, one would expect all areas with large receptive fields to be sensitive. In our data, this is not the case. Prior work has shown that TO1 and TO2, two retinotopic areas with a large degree of correspondence with our functionally defined area MT (Amano et al., 2009), have receptive fields that are slightly larger than those of retinotopic areas LO1 and LO2, and are much larger than those in V3 (Amano et al., 2009; Fig. 6). Importantly, both our functionally defined LOC (which overlap with LO1 and LO2; Larsson and Heeger, 2006) and V3 had a parametric response to symmetry, while MT had no response to symmetry whatsoever in our data. Similarly, receptive field sizes in V3A are at least as large as in V3, if not larger (Smith et al., 2001; Dumoulin and Wandell, 2008), and yet we did not find a parametric response in our V3AB region of interest. This pattern of results makes it unlikely that the differences we see are due to receptive field size alone, but future work should examine the effect that the spatial scale of the lattice and retinal eccentricity have on responses in different visual areas.

The relationship between symmetry perception and spatial scale also intersects with a possible role being played by the periodicity of the wallpapers. The parametric responses we see are clearly driven by rotation symmetry, but it is possible that the periodicity of the patterns interacts with symmetries within the unit lattice, perhaps in different ways across visual areas. Future experiments that vary the size and number of unit lattices in each wallpaper group, and look at individual lattices in isolation, will help to resolve these questions.

In the fields of computer vision and pattern recognition, a principled algorithm for automatic classification of the crystallographic groups has been proposed and validated on real images (Liu et al., 2004a). However, a corresponding model for human perception is missing. To resolve this question, it will be important to probe other wallpaper groups, beyond the four we investigated here. Studies of the wallpaper groups that contain reflection symmetry will help to determine whether there is any difference between symmetry in wallpaper groups and the types of symmetry patterns that have been used in previous experiments (Sasaki et al., 2005). Glide reflection has received very little attention in the literature on symmetry perception, and wallpaper groups can provide a principled way to explore this type of symmetry. Behavioral work has shown that while many of the wallpaper groups are distinguishable by naive observers (Landwehr, 2009), there is significant variability in the perceptual similarity among groups (Clarke et al., 2011b). It will be interesting to use a representational similarity analysis approach (Kriegeskorte et al.,

2008) to explore the extent to which perceptual similarities are mirrored in neural responses.

Notes

Supplemental material for this article is available at <http://purl.stanford.edu/wz510cf6829>. This material consists of the MATLAB code for generating the wallpaper stimulus. This material has not been peer reviewed.

References

- Adelson EH (2001) On seeing stuff: the perception of materials by humans and machines. *Proc SPIE* 4299:1–12. [CrossRef](#)
- Ales JM, Appelbaum LG, Cottareau BR, Norcia AM (2013) The time course of shape discrimination in the human brain. *Neuroimage* 67:77–88. [CrossRef Medline](#)
- Amano K, Wandell BA, Dumoulin SO (2009) Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *J Neurophysiol* 102:2704–2718. [CrossRef Medline](#)
- Arcaro MJ, McMains SA, Singer BD, Kastner S (2009) Retinotopic organization of human ventral visual cortex. *J Neurosci* 29:10638–10652. [CrossRef Medline](#)
- Arcizet F, Jouffrais C, Girard P (2008) Natural textures classification in area V4 of the macaque monkey. *Exp Brain Res* 189:109–120. [CrossRef Medline](#)
- Baayen RH, Davidson DJ, Bates DM (2008) Mixed-effects modeling with crossed random effects for subjects and items. *J Mem Lang* 59:390–412. [CrossRef](#)
- Baillet S, Mosher JC, Leahy RM (2001) Electromagnetic brain mapping. *IEEE Signal Processing Mag* 18:14–30. [CrossRef](#)
- Bakin JS, Nakayama K, Gilbert CD (2000) Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations. *J Neurosci* 20:8188–8198. [Medline](#)
- Bar M, Tootell RB, Schacter DL, Greve DN, Fischl B, Mendola JD, Rosen BR, Dale AM (2001) Cortical mechanisms specific to explicit visual object recognition. *Neuron* 29:529–535. [CrossRef Medline](#)
- Bona S, Herbert A, Toneatto C, Silvanto J, Cattaneo Z (2014) The causal role of the lateral occipital complex in visual mirror symmetry detection and grouping: an fMRI-guided TMS study. *Cortex* 51:46–55. [CrossRef Medline](#)
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207–4221. [Medline](#)
- Brewer AA, Liu J, Wade AR, Wandell BA (2005) Visual field maps and stimulus selectivity in human ventral occipital cortex. *Nat Neurosci* 8:1102–1109. [CrossRef Medline](#)
- Brouwer GJ, Heeger DJ (2009) Decoding and reconstructing color from responses in human visual cortex. *J Neurosci* 29:13992–14003. [CrossRef Medline](#)
- Burkhalter A, Felleman DJ, Newsome WT, Van Essen DC (1986) Anatomical and physiological asymmetries related to visual areas V3 and VP in macaque extrastriate cortex. *Vision Res* 26:63–80. [CrossRef Medline](#)
- Clarke ADF, Halley F, Newell AJ, Griffin LD, Chantler MJ (2011a) Perceptual similarity: a texture challenge. In: *Proceedings of the British Machine Vision Conference* (Hoey J, McKenna S, Trucco E, eds), pp 120.1–120.0. Durham, UK: BMVA Press. [CrossRef](#)
- Clarke ADF, Green PR, Halley F, Chantler MJ (2011b) Similar symmetries: the role of wallpaper groups in perceptual texture similarity. *Symmetry* 3:246–264. [CrossRef](#)
- Cottareau BR, McKee SP, Ales JM, Norcia AM (2011) Disparity-tuned population responses from human visual cortex. *J Neurosci* 31:954–965. [CrossRef Medline](#)
- Cottareau BR, Ales JM, Norcia AM (2012) Increasing the accuracy of electromagnetic inverses using functional area source correlation constraints. *Hum Brain Mapp* 33:2694–2713. [CrossRef Medline](#)
- Cottareau BR, Ales JM, Norcia AM (2014) The evolution of a disparity decision in human visual cortex. *Neuroimage* 92:193–206. [CrossRef Medline](#)
- Cottareau BR, Ales JM, Norcia AM (2015) How to use fMRI functional localizers to improve EEG/MEG source estimation. *J Neurosci Methods* 250:64–73. [CrossRef Medline](#)
- Cox RW (1996) AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29:162–173. [CrossRef Medline](#)
- Dakin SC, Watt RJ (1997) The computation of orientation statistics from visual texture. *Vision Res* 37:3181–3192. [CrossRef Medline](#)
- Deco G, Rolls ET (2004) A Neurodynamical cortical model of visual attention and invariant object recognition. *Vision Res* 44:621–642. [CrossRef Medline](#)
- DeYoe EA, Bandettini P, Neitz J, Miller D, Winans P (1994) Functional magnetic resonance imaging (fMRI) of the human brain. *J Neurosci Methods* 54:171–187. [CrossRef Medline](#)
- DiCarlo JJ, Cox DD (2007) Untangling invariant object recognition. *Trends Cogn Sci* 11:333–341. [CrossRef Medline](#)
- Dong X, Methven TS, Chantler MJ (2014) How well do computational features perceptually rank textures? A comparative evaluation. In: *Proceedings of International Conference on Multimedia Retrieval*, pp 281–288. Glasgow, UK: ACM.
- Dougherty RF, Koch VM, Brewer AA, Fischer B, Modersitzki J, Wandell BA (2003) Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *J Vis* 3(10):1 586–598. [CrossRef](#)
- Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. *Neuroimage* 39:647–660. [CrossRef Medline](#)
- Emrith K, Chantler MJ, Green PR, Maloney LT, Clarke AD (2010) Measuring perceived differences in surface texture due to changes in higher order statistics. *J Opt Soc Am A Opt Image Sci Vis* 27:1232–1244. [CrossRef Medline](#)
- Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. *Nature* 369:525. [CrossRef Medline](#)
- Engel SA, Glover GH, Wandell BA (1997) Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7:181–192. [CrossRef Medline](#)
- Fedorov E (1891) Symmetry in the plane (in Russian). In: *Proceedings of the Imperial St. Petersburg Mineralogical Society*, series 2, Vol 28, pp 245–291.
- Freeman J, Simoncelli EP (2011) Metamers of the ventral stream. *Nat Neurosci* 14:1195–1201. [CrossRef Medline](#)
- Freeman J, Ziemba CM, Heeger DJ, Simoncelli EP, Movshon JA (2013) A functional and perceptual signature of the second visual area in primates. *Nat Neurosci* 16:974–981. [CrossRef Medline](#)
- Galerne B, Gousseau Y, Morel JM (2011) Random phase textures: theory and synthesis. *IEEE Trans Image Process* 20:257–267. [CrossRef Medline](#)
- Gattass R, Sousa AP, Gross CG (1988) Visuotopic organization and extent of V3 and V4 of the macaque. *J Neurosci* 8:1831–1845. [Medline](#)
- Geusebroek J-M, Smeulders AW (2002) A physical explanation for natural image statistics. In: *Proceedings of the 2nd International Workshop on Texture Analysis and Synthesis (Texture 2002)* (Chantler M, ed), pp 47–52. Berlin, New York: Heriott-Watt University.
- Groen II, Ghebreab S, Lamme VA, Scholte HS (2012a) Spatially pooled contrast responses predict neural and perceptual similarity of naturalistic image categories. *PLoS Comput Biol* 8:e1002726. [CrossRef Medline](#)
- Groen II, Ghebreab S, Lamme VA, Scholte HS (2012b) Low-level contrast statistics are diagnostic of invariance of natural textures. *Front Comput Neurosci* 6:34. [CrossRef Medline](#)
- Groen II, Ghebreab S, Prins H, Lamme VA, Scholte HS (2013) From image statistics to scene gist: evoked neural activity reveals transition from low-level natural image structure to scene category. *J Neurosci* 33:18814–18824. [CrossRef Medline](#)
- Hämäläinen MS, Sarvas J (1989) Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data. *IEEE Trans Biomed Eng* 36:165–171. [Medline](#)
- Hämäläinen M, Hari R, Ilmoniemi RJ, Knuutila J, Lounasmaa OV (1993) Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain. *Rev Mod Phys* 65:413. [CrossRef](#)
- Harvey BM, Dumoulin SO (2011) The relationship between cortical magnification factor and population receptive field size in human visual cortex: constancies in cortical architecture. *J Neurosci* 31:13604–13612. [CrossRef Medline](#)
- Hegd  J, Van Essen DC (2007) A comparative study of shape representation in macaque visual areas V2 and V4. *Cereb Cortex* 17:1100–1116. [Medline](#)
- Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. *J Physiol* 148:574–591. [CrossRef Medline](#)
- Huk AC, Dougherty RF, Heeger DJ (2002) Retinotopy and functional subdivision of human areas MT and MST. *J Neurosci* 22:7195–7205. [Medline](#)

- Kass M, Witkin A (1987) Analyzing oriented patterns. *Comput Vision Graph Image Process* 37:362–385. [CrossRef](#)
- Kay KN, Winawer J, Rokem A, Mezer A, Wandell BA (2013) A two-stage cascade model of BOLD responses in human visual cortex. *PLoS Comput Biol* 9:e1003079. [CrossRef](#) [Medline](#)
- Knill DC (1998a) Discrimination of planar surface slant from texture: human and ideal observers compared. *Vision Res* 38:1683–1711. [CrossRef](#) [Medline](#)
- Knill DC (1998b) Ideal observer perturbation analysis reveals human strategies for inferring surface orientation from texture. *Vision Res* 38:2635–2656. [CrossRef](#) [Medline](#)
- Knill DC (1998c) Surface orientation from texture: ideal observers, generic observers and the information content of texture cues. *Vision Res* 38:1655–1682. [CrossRef](#) [Medline](#)
- Kourtzi Z, Kanwisher N (2000) Cortical regions involved in perceiving object shape. *J Neurosci* 20:3310–3318. [Medline](#)
- Kovesi P (2000) Phase congruency: a low-level image invariant. *Psychol Res* 64:136–148. [CrossRef](#) [Medline](#)
- Kriegeskorte N, Mur M, Bandettini P (2008) Representational similarity analysis—connecting the branches of systems neuroscience. *Front Syst Neurosci* 2:4. [CrossRef](#) [Medline](#)
- Kuznetsova A, Brockhoff PB, Christensen RHB (2014) lmerTest: Tests for random and fixed effects for linear mixed effect models (lmer objects of lme4 package). R package version 2.0-11.
- Landwehr K (2009) Camouflaged symmetry. *Perception* 38:1712–1720. [CrossRef](#) [Medline](#)
- Larsson J, Heeger DJ (2006) Two retinotopic visual areas in human lateral occipital cortex. *J Neurosci* 26:13128–13142. [CrossRef](#) [Medline](#)
- Li A, Zaidi Q (2000) Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Res* 40:217–242. [CrossRef](#) [Medline](#)
- Liu Y, Collins RT, Tsin Y (2004a) A computational model for periodic pattern perception based on frieze and wallpaper groups. *IEEE Trans Pattern Anal Mach Intell* 26:354–371. [CrossRef](#) [Medline](#)
- Liu Y, Lin W-C, Hays J (2004b) Near-regular texture analysis and manipulation. *ACM Trans Graph* 23:368–376.
- Liu Y, Hel-Or H, Kaplan CS, Van Gool L (2010) Computational symmetry in computer vision and computer graphics. *Foundations Trends Comput Graph Vis* 5:1–195. [CrossRef](#)
- Miller J, Patterson T, Ulrich R (1998) Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology* 35:99–115. [CrossRef](#) [Medline](#)
- Morrone MC, Burr DC (1988) Feature detection in human vision: a phase-dependent energy model. *Proc R Soc Lond B Biol Sci* 235:221–245. [CrossRef](#) [Medline](#)
- Norcia AM, Candy TR, Pettet MW, Vildavski VY, Tyler CW (2002) Temporal dynamics of the human response to symmetry. *J Vis* 2(2):132–139. [CrossRef](#) [Medline](#)
- Okazawa G, Tajima S, Komatsu H (2015) Image statistics underlying natural texture selectivity of neurons in macaque V4. *Proc Natl Acad Sci U S A* 112:E351–E360. [CrossRef](#) [Medline](#)
- Olshausen BA (2013) Highly overcomplete sparse coding. In: *Proc. SPIE Human Vision and Electronic Imaging XVIII*, p 86510S. Bellingham, WA: International Society for Optics and Photonics.
- Olshausen BA, Field DJ (1997) Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res* 37:3311–3325. [CrossRef](#) [Medline](#)
- Owens RA (1994) Feature-free images. *Pattern Recognit Lett* 15:35–44. [CrossRef](#)
- Pascual-Marqui RD, Michel CM, Lehmann D (1994) Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. *Int J Psychophysiol* 18:49–65. [CrossRef](#)
- Polya G (1924) XII. Über die Analogie der Kristallsymmetrie in der Ebene. *Zeitschrift für Kristallographie-Crystalline Materials* 60:278–282.
- Portilla J, Simoncelli E (2000) A parametric texture model based on joint statistics of complex wavelet coefficients. *Int J Comput Vis* 40:49–70. [CrossRef](#)
- R Core Team (2014) R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. <http://www.R-project.org/>.
- Riesenhuber M, Poggio T (2000) Models of object recognition. *Nat Neurosci* 3:1199–1204. [CrossRef](#) [Medline](#)
- Rosin P, Rammler E (1933) Regularities in the distribution of cement particles. *J Inst Fuel* 7:29–33.
- Sasaki Y, Vanduffel W, Knutsen T, Tyler C, Tootell R (2005) Symmetry activates extrastriate visual cortex in human and nonhuman primates. *Proc Natl Acad Sci U S A* 102:3159–3163. [CrossRef](#) [Medline](#)
- Scholte HS, Ghebreab S, Waldorp L, Smeulders AW, Lamme VA (2009) Brain responses strongly correlate with Weibull image statistics when processing natural images. *J Vis* 9(4):291–15. [CrossRef](#) [Medline](#)
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893. [CrossRef](#) [Medline](#)
- Smith AT, Singh KD, Williams AL, Greenlee MW (2001) Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cereb Cortex* 11:1182–1190. [CrossRef](#) [Medline](#)
- Smith SM (2002) Fast robust automated brain extraction. *Hum Brain Mapp* 17:143–155. [CrossRef](#) [Medline](#)
- Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, Niazy RK, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady JM, Matthews PM (2004) Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23:S208–S219. [CrossRef](#) [Medline](#)
- Thomas OM, Cumming BG, Parker AJ (2002) A specialization for relative disparity in V2. *Nat Neurosci* 5:472–478. [CrossRef](#) [Medline](#)
- Tootell RB, Mendola JD, Hadjikhani NK, Ledden PJ, Liu AK, Reppas JB, Sereno MI, Dale AM (1997) Functional analysis of V3A and related areas in human visual cortex. *J Neurosci* 17:7060–7078. [Medline](#)
- Tootell RB, Hadjikhani N, Hall EK, Marrett S, Vanduffel W, Vaughan JT, Dale AM (1998) The retinotopy of visual spatial attention. *Neuron* 21:1409–1422. [CrossRef](#) [Medline](#)
- Tyler CW, Baseler HA, Kontsevich LL, Likova LT, Wade AR, Wandell BA (2005) Predominantly extra-retinotopic cortical response to pattern symmetry. *Neuroimage* 24:306–314. [CrossRef](#) [Medline](#)
- Vandenbroucke ARE, Fahrenfort JJ, Meuwese JDI, Scholte HS, Lamme VAF (2014) Prior knowledge about objects determines neural color representation in human visual cortex. *Cereb Cortex*. Advance online publication. Retrieved 8 December 2015. [CrossRef](#) [Medline](#)
- Varma M, Zisserman A (2003) Texture classification: are filter banks necessary? In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp II-691–698. Piscataway, NJ: IEEE.
- von der Heydt R, Zhou H, Friedman HS (2000) Representation of stereoscopic edges in monkey visual cortex. *Vision Res* 40:1955–1967. [CrossRef](#) [Medline](#)
- Wade AR, Brewer AA, Rieger JW, Wandell BA (2002) Functional measurements of human ventral occipital cortex: retinotopy and colour. *Philos Trans R Soc Lond B Biol Sci* 357:963–973. [CrossRef](#) [Medline](#)
- Wade D (1993) *Crystal and dragon: the cosmic dance of symmetry and chaos in nature, art, and consciousness*. Rochester, VT: Destiny Books.
- Wandell BA, Winawer J (2011) Imaging retinotopic maps in the human brain. *Vision Res* 51:718–737. [CrossRef](#) [Medline](#)
- Wandell BA, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. *Neuron* 56:366–383. [CrossRef](#) [Medline](#)
- Zhou H, Friedman HS, von der Heydt R (2000) Coding of border ownership in monkey visual cortex. *J Neurosci* 20:6594–6611. [Medline](#)