The HCP 7T Retinotopy Dataset

Noah Benson[1], Keith W. Jamison[2,^], Michael J. Arcaro[3],
An Vu[2,#], Matthew F. Glasser[4,5], Timothy S. Coalson[4], David Van Essen[4],
Essa Yacoub[2], Kamil Ugurbil[2], Jonathan Winawer[1,*], Kendrick Kay[2,*]

[1]Department of Psychology and Center for Neural Science, New York University
[2]Center for Magnetic Resonance Research (CMRR), Department of Radiology, University of Minnesota
[3]Department of Neurobiology, Harvard Medical School
[4]Department of Neuroscience, Washington University in St. Louis
[5]St. Luke's Hospital, St. Louis

*Contributed equally
[^]*Current address:* Department of Radiology, Weill Cornell Medical College
[#]*Current address:* San Francisco VA Health Care System, Center for Imaging of Neurodegenerative Disease

Keywords: retinotopy, visual cortex, fMRI, population receptive fields, topography, parcellation

Abstract

About a quarter of human cerebral cortex is tiled with maps of the visual field. These maps can be measured with functional magnetic resonance imaging (fMRI) while subjects view spatially modulated visual stimuli, also known as ‘retinotopic mapping’. One of the datasets collected by the Human Connectome Project (HCP) involved ultra-high-field (7 Tesla) fMRI retinotopic mapping in 181 healthy adults (1.6-mm resolution), yielding the largest freely available collection of retinotopy data. Here, we describe the experimental paradigm and the results of model-based analysis of the fMRI data. These results provide estimates of population receptive field position and size. Our analyses include both results from individual subjects as well as results obtained by averaging fMRI time-series across subjects at each cortical and subcortical location and then fitting models. Both the group-average and individual-subject results reveal robust signals across much of the brain, including occipital, temporal, parietal, and frontal cortex as well as subcortical areas. The group-average results agree well with previously published parcellations of visual areas. In addition, split-half analyses demonstrate strong within-subject reliability, further evidencing the high quality of the data. We make publicly available the analysis results for individual subjects and the group average, as well as associated stimuli and analysis code. These resources provide an opportunity for studying fine-scale individual variability in cortical and subcortical organization and the properties of high-resolution fMRI. In addition, they provide a measure that can be combined with other HCP measures acquired in these same participants. This enables comparisons across groups, health, and age, and comparison of organization derived from a retinotopic task against that derived from other measurements such as diffusion imaging and resting-state functional connectivity.

Introduction

The central nervous system maps sensory inputs onto topographically organized representations. In the field of vision, researchers have successfully exploited functional magnetic resonance imaging (fMRI) to noninvasively measure visual field representations ('retinotopy') in the living human brain (Engel et al., 1994; Sereno et al., 1995; DeYoe et al., 1996; Engel et al., 1997). These efforts enable parcellation of visual cortex into distinct maps of the visual field, thereby laying the foundation for detailed investigations of the properties of visual cortex (*parcellation references*: Abdollahi et al., 2014; Benson et al., 2014; Wang et al., 2015; *review references*: Tootell et al., 1996; Wandell et al., 2007; Silver and Kastner, 2009; Wandell and Winawer, 2011). One of the datasets acquired by the Human Connectome Project (HCP) (Ugurbil et al., 2013; Van Essen et al., 2013) was a 7T fMRI retinotopy experiment. This experiment, conducted in 181 healthy young adults, involved carefully designed stimuli and a substantial amount of fMRI data (30 minutes, 1,800 time points) acquired at high spatial and temporal resolution (1.6-mm isotropic voxels, 1-second sampling). Although retinotopy is routinely measured in small groups of subjects by individual laboratories in support of various research projects, to date there has not been a large publicly available set of retinotopic measurements.

In this paper, we describe the design of the retinotopy experiment and demonstrate the analyses that we have performed on the fMRI data. We adopt a model-based analysis approach in which a computationally intensive nonlinear optimization is performed to determine parameters of a population receptive field (pRF) model (Dumoulin and Wandell, 2008; Kay et al., 2013; Wandell and Winawer, 2015). The results include estimates of pRF position (angle and eccentricity) and pRF size for each 'grayordinate' (cortical surface vertex or subcortical voxel), and can be used to define retinotopic maps in the brain. We show that the HCP retinotopy data provide high-quality pRF results in many parts of occipital, temporal, parietal, and frontal cortex. We make freely available these pRF results, as well as associated stimuli and analysis code, at an Open Science Framework web site (https://osf.io/bw9ec/). The pRF results are also accessible via the BALSA database (http://balsa.wustl.edu; Van Essen et al., 2017), downloadable as 'scene files' that can be visualized using Connectome Workbench software (see Supplementary Information). The neuroscience community at large can now exploit these resources for a variety of purposes, such as developing normative models, mapping new brain areas, analyzing connectomics, characterizing individual differences, and comparing with other suitably aligned datasets (either published or ongoing).

Results

Here we present a summary of the data quality and example results from the HCP 7T Retinotopy Dataset. The stimuli and analyses are detailed in the Methods and are described here very briefly. Each of 181 subjects participated in six 5-minute pRF mapping runs. The stimuli comprised colorful object textures windowed through slowly moving apertures (Figure 1A). The colorful object textures were used because they produce high signal-to-noise ratio in high-level visual areas. The apertures were clockwise or counterclockwise rotating wedges, expanding or contracting rings, or bars that swept across the visual field in several directions (Figure 1B).

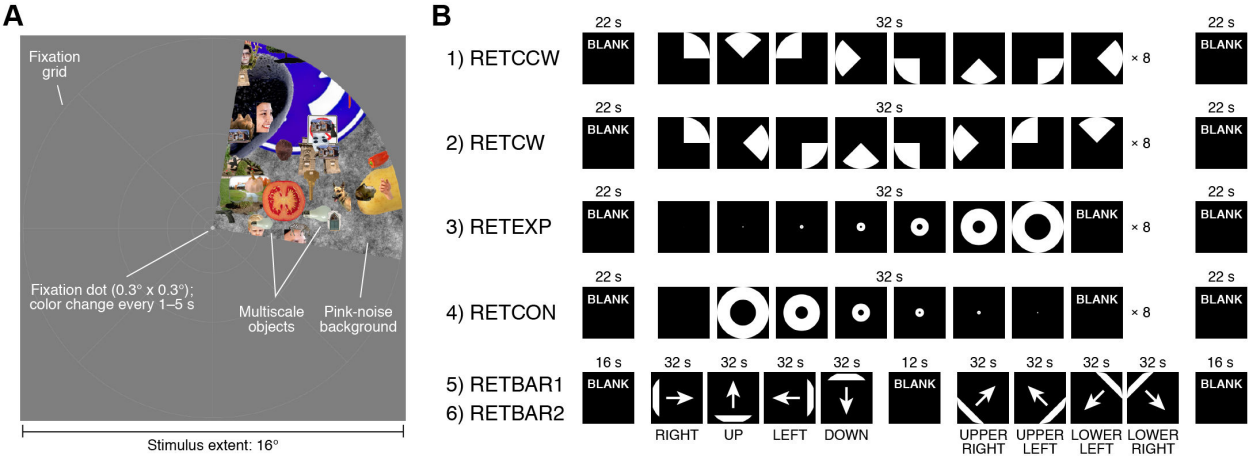

Figure 1. Schematic of experiment. (A) Example stimulus frame. The stimulus consisted of a dynamic colorful texture (composed of objects at multiple scales placed on a pink-noise background) presented within a slowly moving aperture. The aperture and texture were updated at 15 Hz. Subjects were instructed to fixate on a small fixation dot and to press a button whenever its color changed. A fixation grid was provided to aid fixation. (B) Run design. Six 300-s runs were acquired. The temporal structure of the runs is depicted. The first two runs involved a rotating wedge (RETCCW, RETCW), the second two runs involved an expanding or contracting ring (RETEXP, RETCON), and the last two runs involved a moving bar (RETBAR1, RETBAR2).

The resource we provide with this paper is a large set of population receptive field (pRF) model solutions. We define the pRF as the region of the visual field within which a visual stimulus elicits an increase in response from the pooled neural activity reflected in fMRI measurements, and can be summarized by the pRF's angle, eccentricity, and size (Figure 2A). The total dataset consists of 181 individual subjects and 3 group averages. The 3 group averages reflect two split-halves of the subjects as well as all 181 subjects. For each of the 181 individuals and the 3 group averages, we solved 3 sets of models: one from the concatenation of all 6 runs (300 seconds per run, 1,800 time points), one from the first half of each run (150 seconds per run, 900 time points), and one from the second half of each run (150 seconds per run, 900 time points). For each subject or group average and for each of the 3 types of model fits, we obtained model solutions for the 91,282 cortical vertices and subcortical voxels ('grayordinates' spaced on average 2 mm apart). Each model solution yielded 6 numbers: angle, eccentricity, pRF size, gain, variance explained, and mean BOLD signal. Therefore in total, the pRF model solutions that we provide consist of 184 'subjects' (181 individuals plus 3 groups) x 91,282 grayordinates x 3 model fits x 6 quantities (Figure 2B).

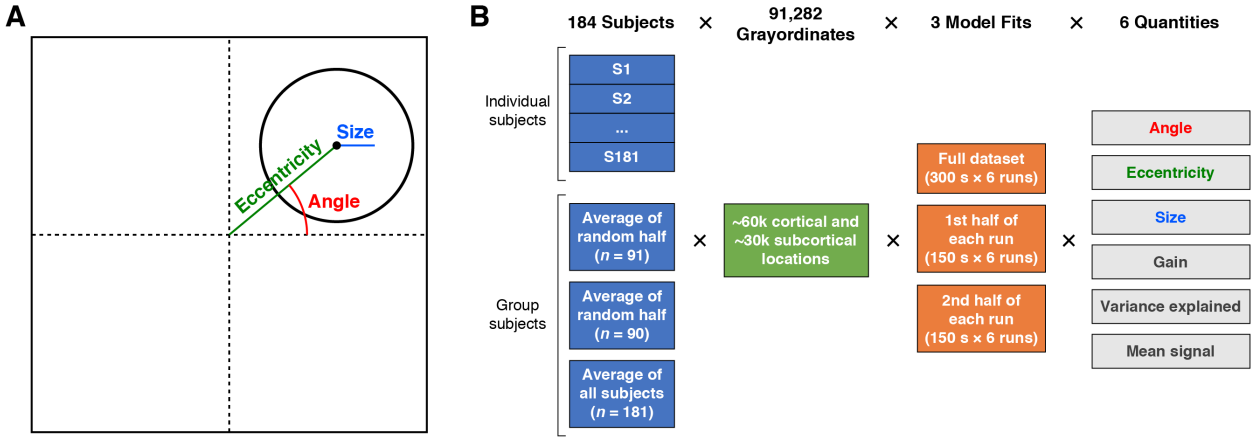

Figure 2. pRF model solutions provided in this resource. (A) pRF parameters. Each pRF is described by a 2D Gaussian. Angle is the rotation of the center of the Gaussian with respect to the positive *x*-axis. Eccentricity is the distance between the center of gaze and the center of the Gaussian. Size is defined as one standard deviation of the Gaussian (the depicted circle is drawn at two standard deviations away from the center). Angle is in units of degrees of polar angle, whereas eccentricity and size are in units of degrees of visual angle. (B) pRF model solutions. We solved pRF models for 181 individual subjects and 3 group-average pseudo-subjects (the average of split-halves of the subjects or of all subjects). For each of the 184 subjects, 3 types of models were fit: one reflecting the complete set of runs and two reflecting split-halves of the runs. Model fits were obtained independently for each of 91,282 grayordinates, yielding 6 quantities. The total dimensions of the pRF model solutions are 184 subjects x 91,282 grayordinates x 3 model fits x 6 quantities.

The particular form of the pRF model we employed assumes that each voxel's pRF is a 2D isotropic Gaussian and that contrast within the pRF is summed sublinearly according to a static power-law nonlinearity with exponent 0.05 (Kay et al., 2013). The sub-additive exponent is used to obtain more accurate pRF solutions, but since it was fixed for all models we do not analyze it further. The pRF size we report is one standard deviation of the Gaussian that describes responses to point stimuli (see Methods).

Group-average results

Cortical data

We first summarize the pRF model solutions from the group-average dataset. Group-average results were obtained by taking the time-series data from individual subjects (aligned using MSMAll to HCP's average cortical surface space fs_LR; see Methods), computing the across-subject average of the time-series data observed at each grayordinate, and then fitting a pRF model to the time-series data at each grayordinate. For visualization, we map the results from fs_LR space to *fsaverage* space and plot the results on the *fsaverage* surface that has been inflated, spherized, and flattened (Figures 3–4). We also provide visualizations of the results on the inflated *fsaverage* surface using dynamic rotating movies (Supplementary Movies 1–12).

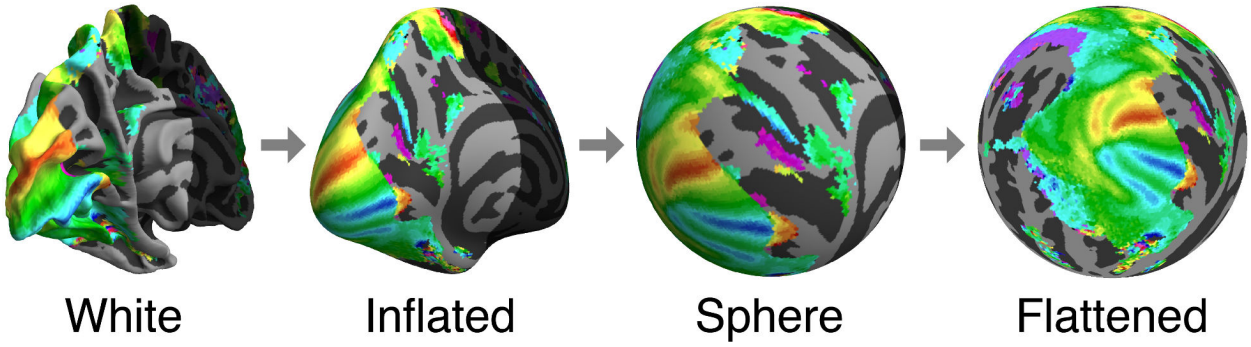

Figure 3. Cortical surface visualization. Cortical surfaces are inflated, warped to a sphere, and rotated with shading discarded to render as an orthographic projection. The regions of the first two surfaces (white, inflated) that are not visible in the final view are darkened. Here we depict the thresholded group-average angle results (see Figure 4) to provide a visual reference across the transformations.

The effect of averaging the time-series data across subjects differs across the cortex, depending on how well pRF parameters match between subjects given the MSMAll alignment. Prior work has shown that the V1–V3 maps have highly regular topography and are well aligned to measures of anatomy, such as surface curvature (Hinds et al., 2008; Benson et al., 2012; 2014) and myelination (Abdollahi et al., 2014), and to measures of function such as resting-state connectivity (Raemaekers et al., 2014; Bock et al., 2015). Therefore, these maps are likely to be well aligned across subjects, and averaging will preserve many of the retinotopic features found in the maps of individual subjects. In particular, the angle and eccentricity maps show clear and expected patterns in V1–V3 (Figure 4, second and third columns), and the variance explained is greater than 75% (cyan regions in the fifth column of Figure 4). As expected, from the lower to upper bank of the calcarine sulcus, there is a smooth progression from the upper vertical meridian through the contralateral horizontal meridian to the lower vertical meridian (blue-cyan-green-yellow-red sweep in the angle colormaps). The angle map reverses at the lips of the calcarine sulcus, with mirror-reversed and approximately quarter-field representations in the bordering dorsal and ventral V2 maps and dorsal and ventral V3 maps. As expected, the eccentricity map is in register across V1–V3, progressing from foveal to peripheral representations from near the occipital pole towards medial and anterior directions (blue-magenta-red-yellow-green progression in the eccentricity colormap). The pRF size map has some of the same features of the eccentricity map, exhibiting smaller sizes near the occipital pole and larger sizes in the mid-peripheral regions of V1–V3. However, in the more peripheral portions of the maps, the size estimates are smaller than expected due to stimulus edge effects (blue rim around the anterior/medial edge of the V1–V3 maps; see Discussion).

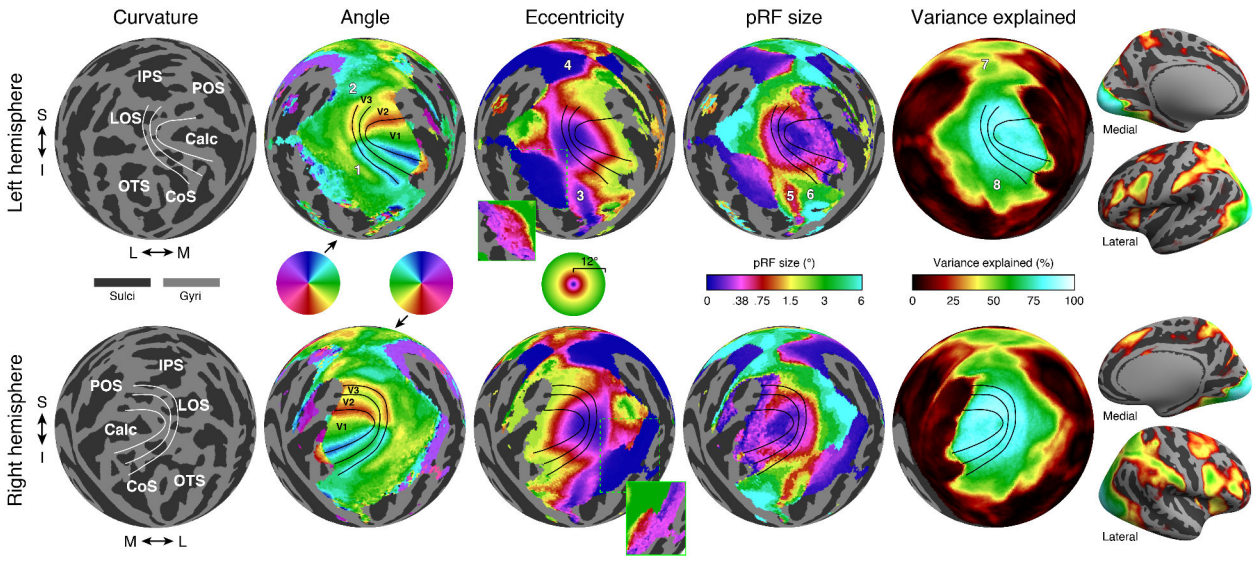

Figure 4. Group-average results. pRF model solutions are mapped from fs_LR space to *fsaverage* using nearest-neighbor interpolation and then visualized (see Methods). The first column shows the *fsaverage* curvature map. White lines are hand-drawn borders of V1, V2, and V3 based on the angle results. Labels indicate several major posterior sulci. The second through fourth columns show angle, eccentricity, and pRF size maps from the group-average dataset (with areal boundaries now shown in black). These maps are thresholded at 9.8% variance explained (see Methods). In the eccentricity maps, the insets marked with green show the same results but with the entire color range corresponding to 0–0.5°—this demonstrates that the large uniform swath of blue in the main figure actually has gradients of near-foveal eccentricities. The fifth column shows variance explained. Finally, the images on the right show thresholded variance explained on inflated left and right surfaces, demonstrating the existence of robust signals in other parts of cortex. Labels: S = superior, I = inferior, M = medial, L = lateral, IPS = intraparietal sulcus, LOS = lateral occipital sulcus, POS = parieto-occipital sulcus, Calc = calcarine sulcus, OTS = occipitotemporal sulcus, CoS = collateral sulcus.

In cortical locations where retinotopic parameters are variable across subjects (even after registration using MSMAll), the group-average results will preserve less of the detail from individual subjects. Nonetheless, there is a large amount of structure in the group-average results beyond V1–V3, and some clear patterns are evident. The angle maps show the expected progression from upper to lower field ventral to V3, and from lower to upper field dorsal to V3 (locations 1 and 2 in Figure 4), consistent with measurements of ventral (McKeefry et al., 1997; Kastner et al., 2001; Wade et al., 2002) and dorsal (Tootell et al., 1997; Press et al., 2001) occipital cortex. The eccentricity map also shows clear large-scale organization throughout large expanses of parietal and temporal cortex. One feature of the eccentricity maps is multiple distinct, foveal representations: in addition to the foveal representation of V1–V3 at the occipital pole, the eccentricity maps show distinct foveal representations in ventral temporal cortex and parietal cortex (locations 3 and 4 in Figure 4), consistent with many prior studies (Tootell et al., 1997; Wade et al., 2002; Wandell et al., 2005; Swisher et al., 2007). Near both of these distinct foveal representations, there are foveal to peripheral gradients along the lateral to medial direction.

The pRF size map also shows a variety of large-scale patterns. In ventral temporal cortex, there is a small-to-large size gradient from the fusiform gyrus to the collateral sulcus (locations 5 and 6 in Figure 4).

These regions roughly correspond to the locations of face-selective and place-selective cortex, respectively (Kanwisher et al., 1997; Epstein and Kanwisher, 1998; Malach et al., 2002; Grill-Spector and Weiner, 2014). More generally, pRF sizes tend to be larger outside V1–V3, as expected from both single-unit and fMRI measurements (Maunsell and Van Essen, 1983; Tootell et al., 1997; Smith et al., 2001; Dumoulin and Wandell, 2008). Finally, the variance explained map shows that robust signals can be found not only within V1–V3 but also in higher-level areas. Variance explained tends to be above 50% in regions that are ventral, lateral, and dorsal to the V1–V3 maps, including much of ventral temporal cortex and the intraparietal sulcus (locations 7 and 8 in Figure 4). Furthermore, for nearly all cortical locations that survive the variance explained threshold, pRF model parameters are highly reliable. This can be verified by comparing model fits obtained for group-average subject 184 across split-halves of each run and by comparing model fits for the full dataset across group-average split-half subjects 182 and 183 (results not shown; pRF model solutions available online).

Relationship to cortical parcellations

Many features of the group-average results are in good agreement with published parcellations of visual areas, particularly near the posterior occipital pole (Figure 5). The Wang et al. maximum probability atlas (Wang et al., 2015) includes 25 regions of interest (ROIs) per hemisphere. Ten of these in posterior cortex are clearly aligned with expected features of the polar angle maps: V1v/V1d, V2v/V2d, V3v/V3d, V3A, V3B, LO-1, and hV4. In each of these 10 ROIs, one or more borders lie on an angle reversal. For example, the V1d/V2d border lies on a lower-field angle reversal, and the V1v/V2v border lies on an upper-field angle reversal. Other maps such as LO-2, TO-1/2, IPS maps, VO-1/2, and PHC-1/2 show contralateral representations but no clear progressions of angle, in part due to blurring from group averaging (i.e., averaging cortical locations from different subjects that represent different portions of the visual field). Similarly, for the Glasser et al. (Glasser et al., 2016) parcellation, several ROIs are well aligned with features of the retinotopic maps, particularly V1, V2, V3, V4, and V3A. In several map clusters in the Wang et al. atlas, there are clear eccentricity gradients: the IPS0–2 maps show a clear foveal-to-peripheral gradient along the medial-to-lateral direction, as do the V1–V3 maps and the VO-1/2 maps. In the Glasser et al. atlas, several regions fall within iso-eccentricity regions. For example, the PH and TE2p ROIs are clearly foveal, whereas the adjacent FFC ROI is more peripheral, with the border between these ROIs corresponding to a sharp change in eccentricity.

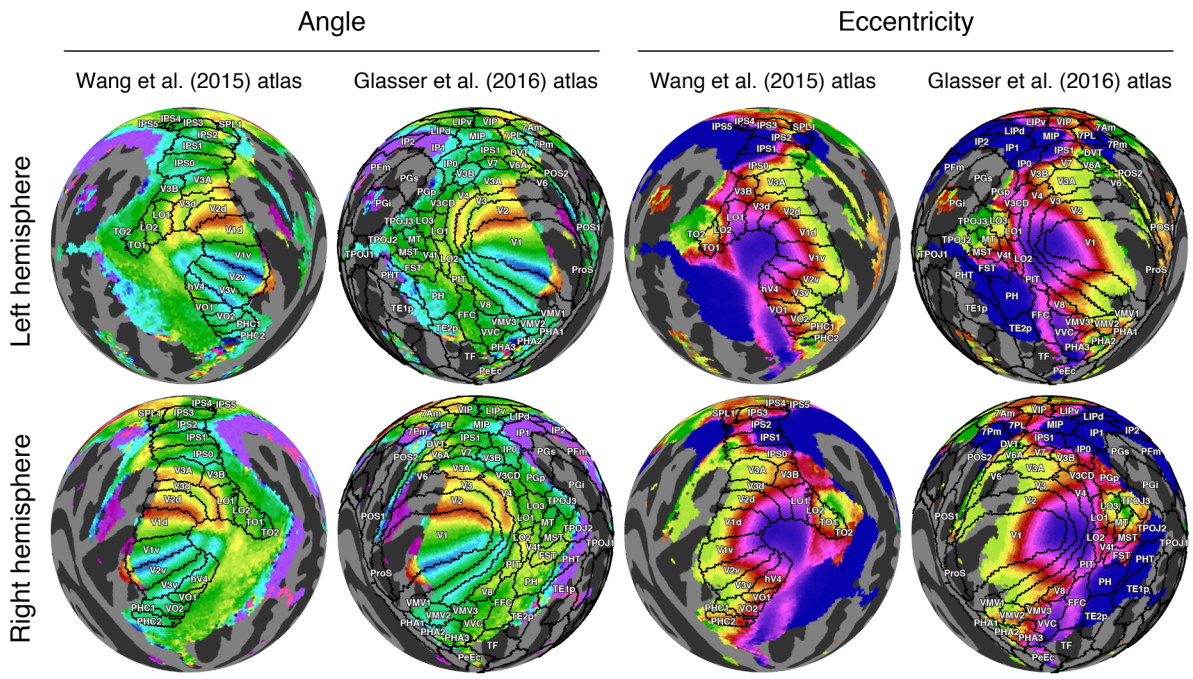

Figure 5. Relationship between group-average results and cortical parcellations. The angle and eccentricity maps from the group-average dataset are re-plotted from Figure 4. Superimposed on the maps are regions of interest from the maximum probability atlas of Wang et al. (Wang et al., 2015) and cortical parcellations from Glasser et al. (Glasser et al., 2016).

Subcortical data

The HCP 7T Retinotopy Dataset includes subcortical results in addition to cortical results. The subcortical fMRI data were aligned using FNIRT nonlinear volume registration based on T1-weighted image intensities (Glasser et al., 2013). Several subcortical nuclei have retinotopic maps that have been previously measured using fMRI (Schneider et al., 2004; Schneider and Kastner, 2005; Cotton and Smith, 2007; Katyal et al., 2010; Arcaro et al., 2015; DeSimone et al., 2015). In contrast to cortex, subcortical structures are not easily represented as 2D surfaces, and hence it is more difficult to visualize complete maps. Nonetheless, slices through subcortical structures reveal clear, high-quality pRF model solutions in the group-average dataset (Figure 6). In particular, we see expected structure in visual nuclei such as the lateral geniculate nucleus (LGN), superior colliculus (SC), and ventral pulvinar (vPul1/2). Within these regions, there are clear representations of the contralateral visual field. As expected, the visual field maps of the LGN and pulvinar are both inverted with smooth progressions from the upper visual field located ventrally to the lower visual field located dorsally. In the superior colliculus, there is a smooth progression from the upper visual field (anterior and medial) to the lower visual field (posterior and lateral).

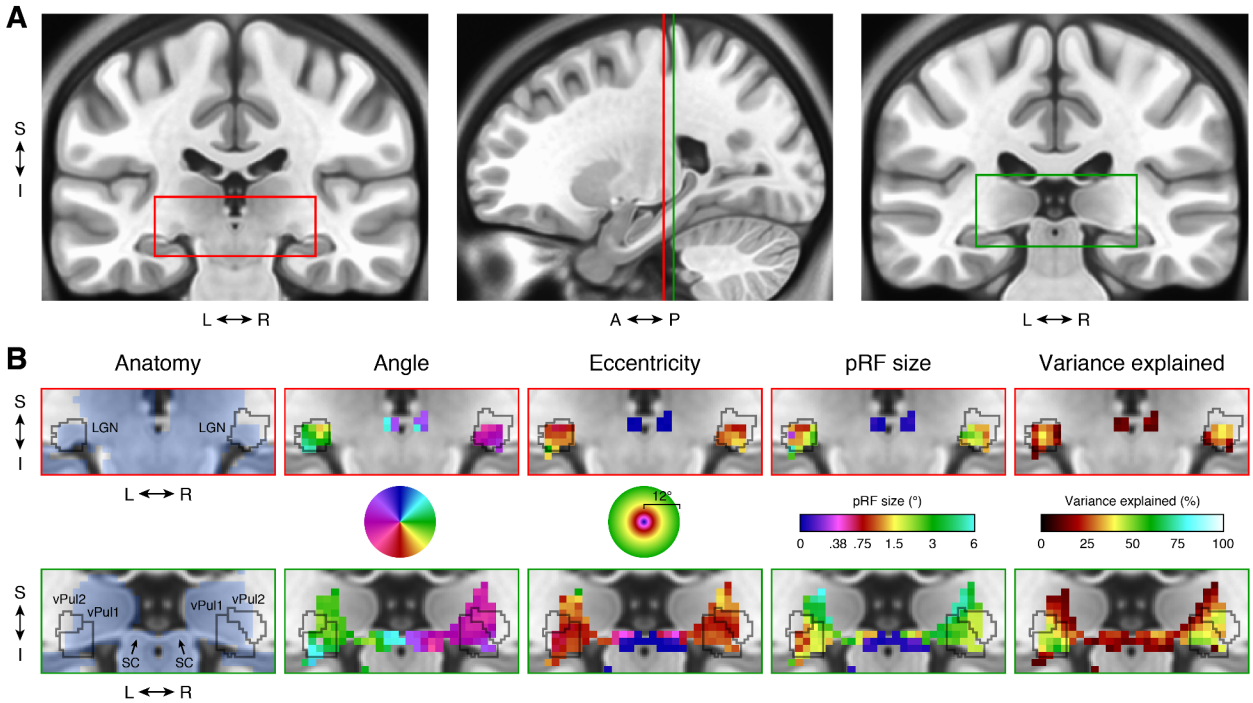

Figure 6. Subcortical results. (A) Anatomical location. The two coronal slices (y = 25, far left; y = 30, far right) show the MNI average anatomy (ICBM 152 nonlinear symmetric atlas 2009b, 0.5-mm resolution). The red and green rectangles mark the regions detailed in panel B. Vertical lines on the sagittal slice (x = 23) indicate the locations of the two coronal slices. (B) pRF results. The upper row highlights the left and right LGN and the lower row highlights the pulvinar and superior colliculus. Outlines of the LGN and ventral pulvinar (vPul1/2) are taken from Arcaro et al. (2015). All pRF results are from the group-average dataset and are thresholded at 9.8% variance explained, as in Figure 4. Colormaps are identical to those in Figure 4 except that only the left-hemisphere angle colormap is used. The blue shading in the anatomy column indicates voxels that are included in the CIFTI subcortical data mask.

Individual-subject results

In addition to group-average results, we also computed pRF model solutions for the 181 individual subjects. We summarize results in several ways, including quantifying the amount of variance explained by the pRF model, inspecting maps in individual subjects, and assessing within-subject reliability of pRF parameters. These analyses reveal that overall data quality is high.

Variance explained

We quantified variance explained within atlas-defined ROIs. We defined one ROI as the union of the 50 maps found in the Wang et al. maximum probability atlas (25 maps per hemisphere) and a second ROI as the union of the V1–V3 maps from the same atlas (Wang et al., 2015). The V1–V3 ROI is a subset of the larger ROI. Because these ROIs are defined based on group-average anatomy, they do not necessarily conform to each individual subject's retinotopic maps, but they provide a simple objective method for region definition. Within the union of the 50 maps, we computed for each subject the median variance explained across grayordinates, yielding one number per subject. The median of this number

across the 181 subjects was 17% (Figure 7A). Within just the V1–V3 maps, the median of the median variance explained was substantially higher, at 44%. For comparison, we estimate that for grayordinates not sensitive to the experimental paradigm, the variance explained by the pRF model is less than 1%. This can be seen by inspecting the large peak in the histogram of variance explained across all grayordinates from all individual subjects (Figure 7C).

Cortical maps

For map visualization, we selected three representative subjects: the subjects at the 25th, 50th, and 75th percentiles with respect to median variance explained across regions in the Wang et al. atlas (see red lines in Figure 7A). For simplicity we show only the left hemisphere, and we re-plot the group-average results for comparison. The three depicted subjects have clear retinotopic maps in occipital cortex, as seen in the angle and eccentricity results (Figure 7C). In each subject, the angle maps reveal the boundaries of V1–V3, and the eccentricity maps are in register across visual field maps around the occipital pole. The locations of the V1–V3 boundaries differ slightly across the subjects, as seen by comparing the angle reversals and the V1–V3 boundary lines that were drawn based on the group-average results. This suggests that even after alignment using state-of-the-art algorithms guided by folding and areal features (MSMAll), there may still be some differences in retinotopic structure. This is not surprising given the more than two-fold individual variability in the sizes of V1 (Amunts et al., 2000) and early extrastriate areas (Dougherty et al., 2003).

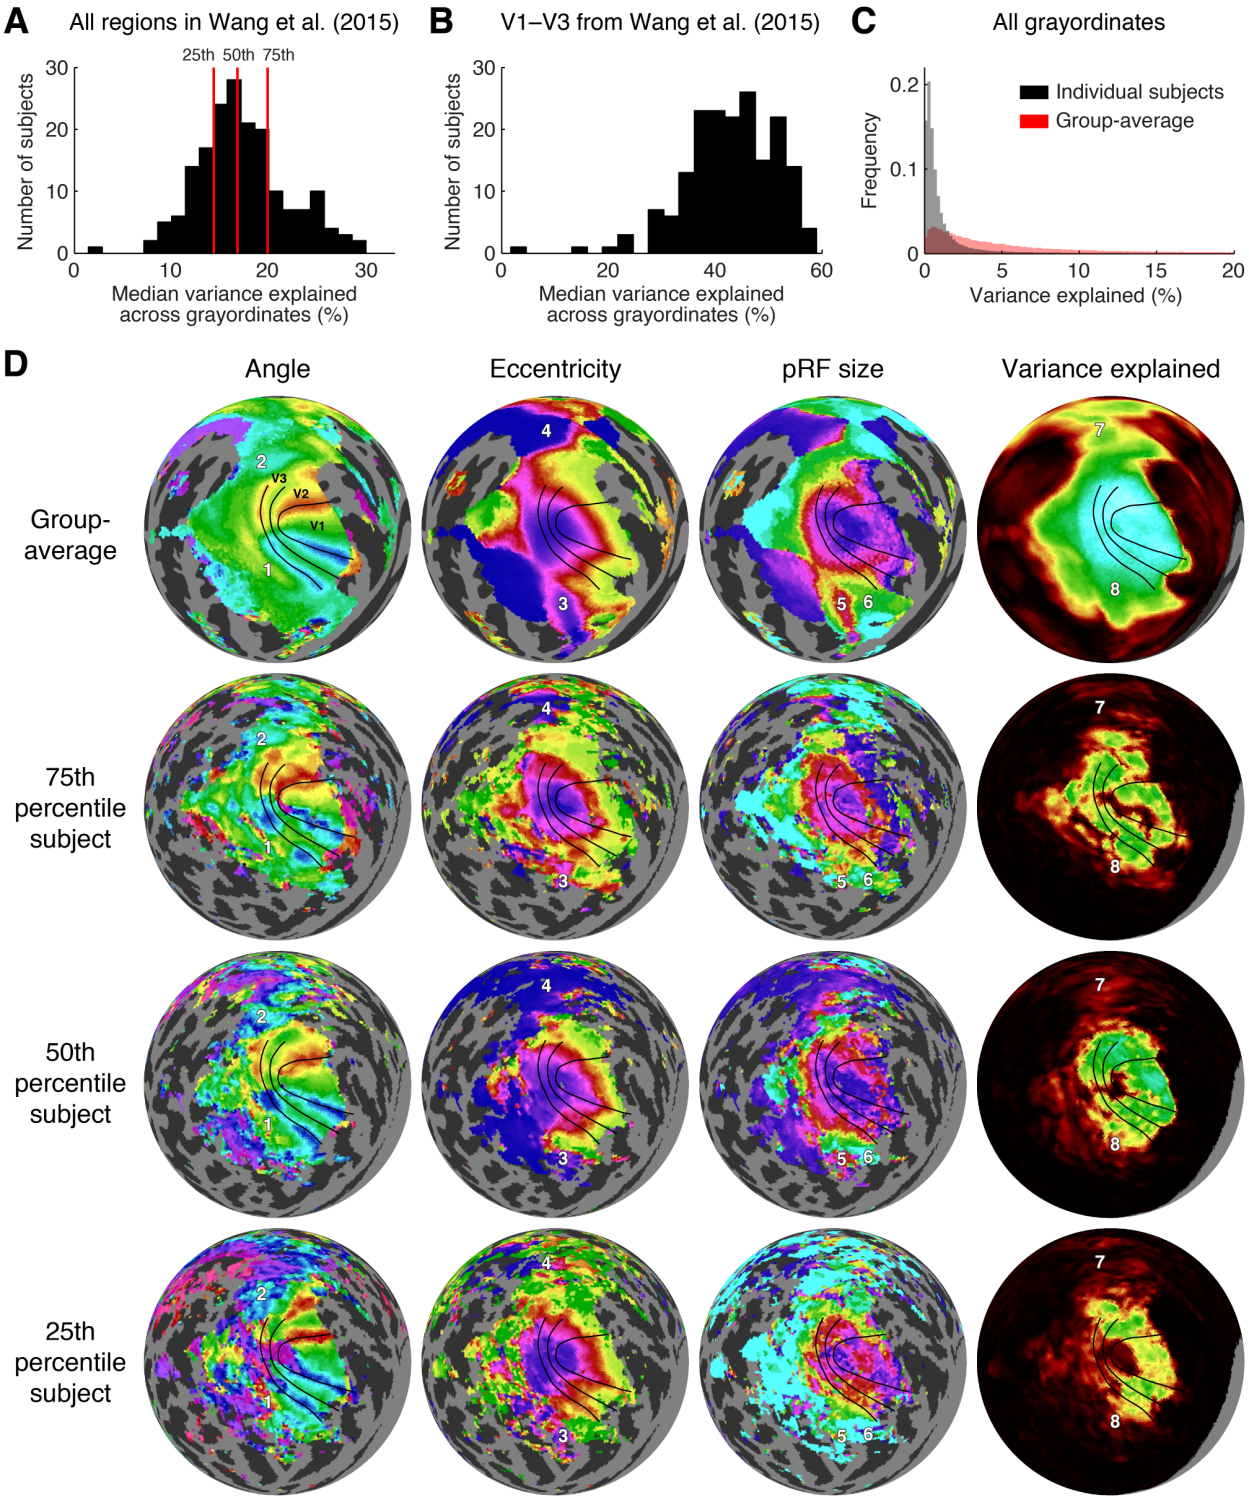

Figure 7. Individual-subject results. (A) Variance explained within all regions of the Wang et al. maximum probability atlas (Wang et al., 2015). For each subject, we computed the median variance explained across grayordinates located within the union of all regions in both hemispheres of the Wang et al. atlas. The histogram shows the distribution of this value across the 181 subjects. The subjects at the 25th, 50th, and 75th

percentiles are indicated by red lines. (B) Variance explained within V1–V3 of the Wang et al. atlas. (C) Histogram of variance explained across all grayordinates in individual subjects and in the group-average subject (bin size 0.2%; histogram counts normalized to sum to 1 in each plot). (D) Maps of pRF parameters (left hemisphere). We re-plot the group-average results (see Figure 4) and show results for the 3 subjects indicated in panel A (HCP subject IDs 164131, 115017, and 536647, corresponding to the 25th, 50th, and 75th percentiles, respectively). Angle, eccentricity, pRF size, and variance explained are plotted as in Figure 4, except that the variance explained threshold used for individual subjects is 2.2% (see Methods). For reference, we show on each map the same V1–V3 boundary lines determined from group-average results in Figure 4.

Beyond V1–V3, several of the features we noted in the group-average results are also generally evident in the individual subjects. For example, the angle maps show a lower-field representation ventral to V3 and an upper-field representation dorsal to V3 (locations 1 and 2 in Figure 7). There are also distinct foveal representations in parietal and temporal cortex (locations 3 and 4), and pRF size gradients in ventral cortex (locations 5 and 6). Because variance explained is generally lower for individual subjects compared to the group average, there are some regions in which the group average may provide useful information that is absent in individual subjects (e.g. location 8). By visual inspection, the overall map quality appears comparable across the three subjects. Since these subjects span the central 50% of variance explained (as detailed previously), this suggests that most of the subjects in the HCP 7T Retinotopy Dataset have good data quality. Additional aspects of individual variability can be readily inspected by scrolling through polar angle and eccentricity maps for all 181 individual subjects in the downloadable Connectome Workbench 'scene' files (see Supplementary Information).

Within-subject reliability

To quantify reliability of pRF parameters for individual subjects, we compared parameter estimates across split-halves of the data. We binned cortical grayordinates into 4 large ROIs which comprise distinct subsets of the regions in the Wang et al. atlas (Wang et al., 2015): posterior (V1–V3), dorsal (V3A/B, IPS0–5), lateral (LO-1/2, TO-1/2), and ventral (VO-1/2, PHC-1/2). We then aggregated grayordinates within each of these ROIs across subjects, and computed 2D histograms comparing parameter estimates across the two model fits (first half of each run; second half of each run).

Angle estimates were highly reliable across splits for all 4 ROIs, indicated by the high density along the diagonals (Figure 8, top row). In addition to demonstrating within-subject reliability, these histograms highlight the fact that angles near the vertical meridian (90º and 270º) are less represented than other angles, an effect observed in many prior studies (Silver et al., 2005; Larsson and Heeger, 2006; Kastner et al., 2007; Swisher et al., 2007; Arcaro et al., 2009; Mackey et al., 2017). The eccentricity histograms (Figure 8, middle row) also show a high degree of reliability, with density highest on and near the diagonal in all 4 ROIs. Note that the dorsal ROI, while reliable, is more foveally biased than other ROIs. Nonetheless, as indicated in both the maps (Figures 4 and 7) and the reliability plots (Figure 8), the dorsal regions contain eccentricities spanning 0 to 8 degrees. Finally, the size estimates were also fairly reliable, though less so than the angle and eccentricity estimates. In agreement with the maps (Figures 4 and 7), posterior maps generally contain the smallest pRFs, with few pRF sizes larger than 3 degrees.

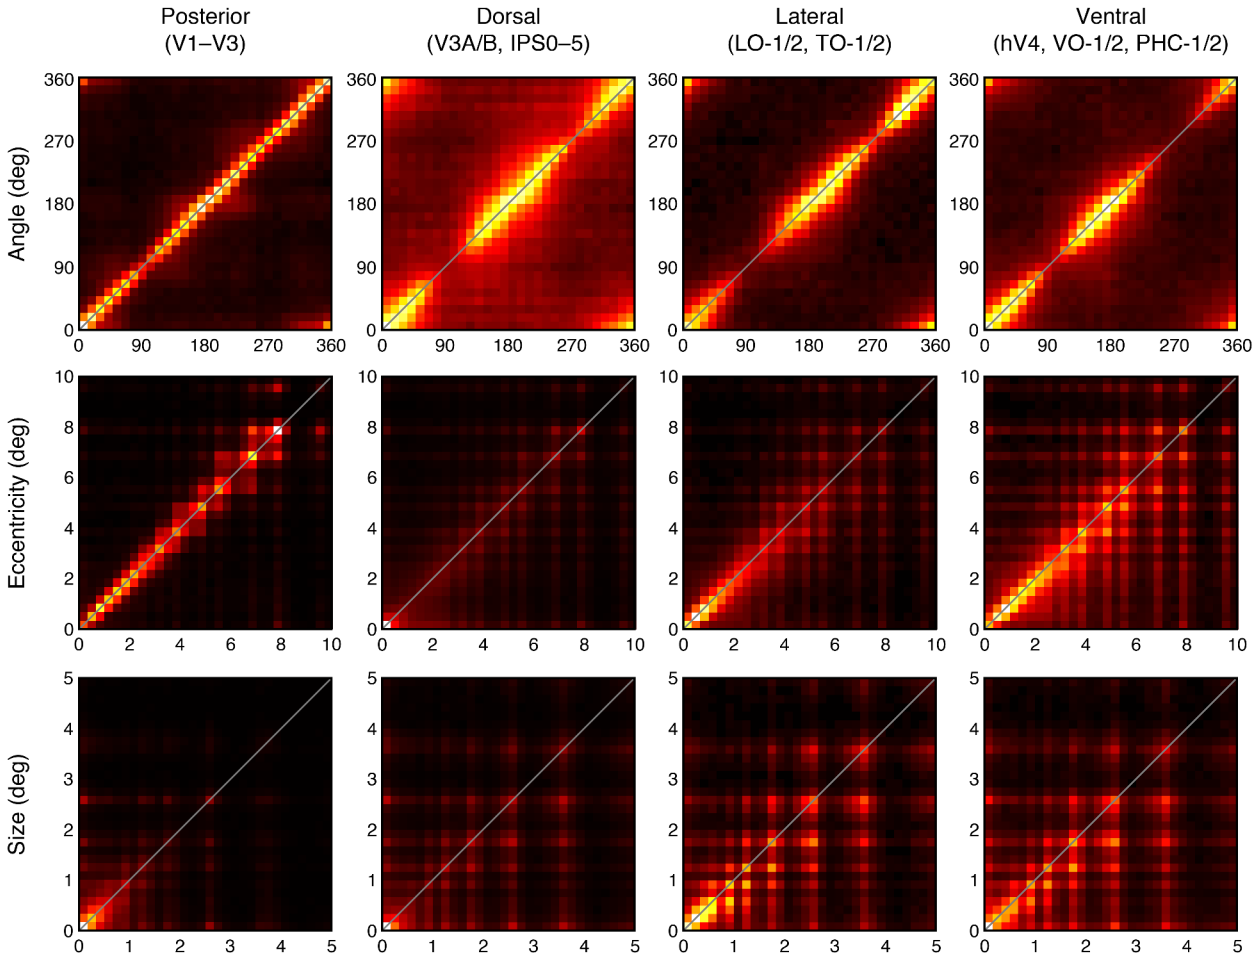

Figure 8. Within-subject reliability of pRF estimates. Estimates of pRF parameters were obtained for two independent splits of the data (first half of each run; second half of each run). Here, we aggregate results across all 181 individual subjects and plot 2D histograms comparing pRF parameter estimates across the two splits of the data (*x*-axis: first split; *y*-axis: second split). To enhance visibility, the colormap was applied to the square root of frequency counts observed in each plot. The Wang et al. atlas (Wang et al., 2015) was used to bin grayordinates into different ROIs (posterior, dorsal, lateral, ventral).

Discussion

We have described the HCP 7T Retinotopy Dataset and the results of fitting pRF models to all 181 individual subjects as well as the group average. To facilitate quantification of model reliability, all datasets were analyzed using split-halves in addition to the full dataset. In addition to the pRF model solutions, we also make available the stimuli and analysis code used to solve the models. This allows the research community to reproduce our analyses and/or re-analyze the time-series data using different techniques. The analyses we conducted are computationally intensive, involving three independent nonlinear optimizations for each grayordinate time series. This resulted in approximately 50 million model fits that necessitated the use of a large-scale compute cluster. By providing the parameters of the solved models, we substantially lower the barrier to entry for scientists to make use of the dataset.

Size and quality of the HCP 7T Retinotopy Dataset

Although researchers frequently collect, and occasionally make public, retinotopy datasets, such datasets have generally included no more than 20 subjects (e.g., Benson et al., 2012). To the best of our knowledge, the HCP 7T Retinotopy Dataset is the largest publicly available dataset by an order of magnitude. In addition to containing many subjects, retinotopic maps are derived from six fMRI runs (a total of 30 minutes of data), making this dataset large both in terms of number of subjects as well as amount of data per subject. Finally, the data were acquired at ultra-high magnetic field strength (7T), providing enhanced signal-to-noise ratio and high spatial and temporal resolution (1.6-mm isotropic voxels, 1-second sampling). The advantages of the dataset are clear. In individual subjects, there are reliable results, even beyond striate and extrastriate cortex. At the group level, the massive averaging of subjects reveals signals in regions of cortex (such as the inferior frontal sulcus) where conventional datasets typically have low signal and are therefore difficult to distinguish from noise.

Limitations of the dataset

Though the dataset has clear value, it is important to understand its limitations and take these limitations into account when interpreting the results. There are several technical issues; we mention a few here, but refer to the Methods for a fuller description. The stimulus size extended to an eccentricity of 8 degrees of visual angle, and so representations of the far periphery are not well measured. Because of cortical magnification of the central visual field, robust signals are found in about half of the surface area of V1 (see Figure 5). Edge effects arise for voxels whose pRF centers are near the limit of the stimulus extent: these voxels are likely to have underestimates of pRF size and a displaced pRF center. Model solutions are somewhat discretized, reflecting the influence of the first-stage grid fit. Model solutions were constrained to have a non-negative gain factor; this may not be appropriate for studying brain regions that exhibit BOLD signal decreases in response to visual stimulation. Finally, we highlight the fact that all analyses have assumptions, and the pRF model solutions we have derived are intrinsically tied to the choice of pRF model. Exploring other models and carefully evaluating model accuracy (Kay et al., 2013) may be important when investigating brain regions that are not well understood.

The neuroscientific interpretation of the pRF results must also be done carefully. Whereas in visual cortex, there is clear interpretation of pRF models in terms of visually responsive population receptive fields, in other parts of the brain, it may be possible to obtain good pRF fits but for different reasons. For example, it is possible that a cortical region indexing cognitive difficulty exhibits response increases when the stimulus is near the fovea because at these points in the experiment, the stimulus is likely to interfere

with the fixation task performed by the subjects. In such a case, the existence of a pRF model solution does not indicate visually driven activity in the conventional sense.

What could the HCP 7T Retinotopy Dataset be used for?

This rich dataset has a wide range of uses. It provides the basis for further analysis of other HCP data; for example, the pRF solutions for an individual subject can be used to determine visual ROIs that could then be used to analyze or validate other HCP measures. Some example applications include the following: (1) The retinotopy dataset can be used for comparison with the HCP's multimodal parcellation (Glasser et al., 2016) (see Figure 5). We have shown that the group-average results approximately agrees with portions of the parcellation, but we did not compare individual-subject results to the parcellation. (2) Identifying visual ROIs and pRF properties within the ROIs can be used in conjunction with resting-state data (Van Essen et al., 2012) to test hypotheses about how maps relate to functional connectivity. (3) The pRF model solutions can be used in conjunction with the working memory dataset (Barch et al., 2013) to study the role of visual cortex in working memory. Many more such applications (e.g. combining retinotopy with the 7T movie data) are possible.

The visuotopic mapping in the Glasser et al. (Glasser et al., 2016) parcellation was based on resting-state fMRI correlations measured across the entire visual field representation. This enabled mapping the full extent of visuotopic areas, but does not provide explicit measurement of specific eccentricities or polar angles within each map. Hence, the current 7T retinotopic maps and the visuotopic organization derivable from resting-state data represent complementary and potentially synergistic information.

The dataset also has a great deal of standalone value, owing to the very large number of subjects. Any examination of the relationship between anatomy and function benefits from having many subjects to characterize the extent of intersubject structure-function variability in an anatomically-normalized format. Averaging retinotopic time-series data across a large number of subjects has revealed that large swaths of cortex not typically studied by vision scientists show evidence of retinotopic organization (see Figure 4); many of these regions would not have clear signals in smaller sample sizes. Many disagreements about the organization of retinotopic maps have remained unresolved due in part to limited numbers of subjects and limited data quality; this dataset may help overcome these limits.

Conclusion

The visual system is one of the primary model systems of human neuroscience, and the resources provided in this paper represent an important step towards more fully characterizing its fundamental organization. The authors believe that the present measurements fill a critical role, both for answering novel scientific questions and for establishing baselines and hypotheses for new experiments. To this end, we have put effort into making all data and analyses fully public and well-documented, and we hope that other researchers will find this dataset enlightening and useful.

Acknowledgements

This work was supported by the Human Connectome Project (1U54MH091657) from the 16 Institutes and Centers of the National Institutes of Health that support the NIH Blueprint for Neuroscience Research; Biotechnology Research Center (BTRC) grant P41 EB015894 from NIBIB; NINDS Institutional Center Core grant P30 NS076408; and R01 EY027401 (J.W.).

Author Contributions

D.V.E. and K.U. planned the experiment. A.V., E.Y., and K.U. developed and optimized acquisition sequences and protocols. A.V. and K.K. designed the experiment. N.B., K.J., M.A., M.G., and K.K. analyzed the data. T.C., M.G., and D.V.E. prepared results for the BALSA database. N.B., J.W., and K.K. wrote the paper. All authors discussed the results and edited the manuscript.

Methods

Subjects

Complete retinotopy datasets (six fMRI runs) were acquired for a total of 181 subjects (109 females, 72 males), age 22–35, as part of the Young Adult Human Connectome Project (HCP) (https://www.humanconnectome.org/study/hcp-young-adult). All subjects also participated in ~4 hours of multimodal MRI data acquisition on a customized Siemens 3T 'Connectom' scanner at Washington University in St. Louis (Van Essen et al., 2013) as well as extensive behavioral and demographic assessments (Barch et al., 2013). All subjects had normal or corrected-to-normal visual acuity. The subjects include 53 pairs of genetically confirmed identical twins (106 individuals), 34 pairs of fraternal twins (68 individuals), 2 pairs of non-twin siblings (4 individuals), and 3 individuals whose twins/siblings were not included. For genetic details, researchers must apply for access to "Restricted Data" on ConnectomeDB.

Structural image acquisition and pre-preprocessing

T1-weighted (T1w) and T2-weighted (T2w) structural scans at 0.7-mm isotropic resolution were acquired at 3T and used as the anatomical substrate for the retinotopy data. White and pial cortical surfaces were reconstructed from the structural scans using the HCP Pipelines (Glasser et al., 2013). Surfaces were aligned across subjects to the HCP 32k fs_LR standard surface space using first a gentle folding-based registration 'MSMSulc' and then a more aggressive areal-feature-based registration 'MSMAll' that was driven by myelin maps, resting-state network maps, and 3T resting-state visuotopic maps (Robinson et al., 2014; Glasser et al., 2016; Robinson et al., 2018). Myelin maps were based on the ratio of T1w/T2w images (Glasser and Van Essen, 2011), normalized using a surface-based atlas to estimate B1+ transmit effects (Glasser et al., 2013). Subcortical volume data were aligned to MNI space using FNIRT nonlinear volume-based registration based on T1w image intensities (Glasser et al., 2013).

fMRI acquisition and pre-processing

Full details on data acquisition and pre-processing are provided elsewhere (Glasser et al., 2013; Vu et al., 2016). In brief, fMRI data were collected at the Center for Magnetic Resonance Research at the University of Minnesota using a Siemens 7T Magnetom actively shielded scanner and a 32-channel receiver coil array with a single-channel transmit coil (Nova Medical, Wilmington, MA). Whole-brain fMRI data were collected at a resolution of 1.6-mm isotropic and 1-s TR (multiband acceleration 5, in-plane acceleration 2, 85 slices). The data were processed using HCP pipelines (Glasser et al., 2013), correcting for head motion and EPI spatial distortion and bringing the fMRI data into alignment with the HCP standard surface space as described above. The data produced by the pipeline are in the CIFTI format, which consists of 91,282 grayordinates that cover both cortical and subcortical brain regions with approximately 2-mm spatial resolution. The fMRI data were also denoised for spatially specific structured noise using multi-run sICA+FIX (Griffanti et al., 2014; Salimi-Khorshidi et al., 2014; Glasser et al., 2017). Differences in slice timing were not corrected since the fast multiband acquisition makes such corrections less important (though slices may differ by as much as 1 s). The dimensions of the pre-processed data are 181 subjects x 91,282 grayordinates x 6 runs x 300 time points. These pre-processed data are available from ConnectomeDB (https://db.humanconnectome.org/).

Note that the fMRI data used in this paper reflect the correct phase-encode directions in the EPI undistortion procedure, unlike an early pre-2018 release of the data. Also, we point out that higher-

resolution CIFTI outputs are also available, consisting of 170,494 grayordinates with approximately 1.6-mm spatial resolution. Only the 2-mm CIFTI outputs were used in this paper.

Stimuli

Retinotopic mapping stimuli were constructed by creating slowly moving apertures and placing a dynamic colorful texture within the apertures. Apertures and textures were generated at a resolution of 768 pixels x 768 pixels, and were constrained to a circular region with diameter 16.0°. The display was uniform gray beyond the circular region.

Texture design

To elicit strong neural responses in high-level visual areas (while also driving responses in early visual areas), we designed a texture composed of colorful visual objects. The texture was constructed by taking objects from Kriegeskorte et al. (2008), preparing these objects at multiple scales, and placing the objects on an achromatic pink-noise background. One hundred (100) distinct texture images were generated. To generate a texture image, we first created an achromatic pink-noise ($1/f$ amplitude spectrum) background. Then, starting at the largest scale and proceeding to smaller scales, objects were randomly selected and placed at random positions in the image, potentially occluding objects already placed, similar to a 'dead leaves' tessellation. There were seven different scales. The object sizes associated with the seven scales were 350, 247, 175, 124, 88, 62, and 44 pixels (decreasing by a factor of $\sqrt{2}$), and the numbers of objects at each scale were 1, 2, 4, 8, 16, 32, and 64 (increasing by a factor of 2). Textures were updated at a rate of 15 Hz (details below). Pilot experiments confirmed that compared to conventional checkerboard patterns, the object texture produces larger BOLD responses and improves test-retest reliability of retinotopic estimates.

Aperture design

The experiment consisted of six runs in which three different types of apertures were presented (wedges, rings, bars). Apertures moved slowly across the visual field, and were occasionally interrupted by blank periods in order to help distinguish between non-visual responses and responses from neurons with very large receptive fields (Dumoulin and Wandell, 2008). Each run lasted 300.0 s. The order of runs was RETCCW, RETCW, RETEXP, RETCON, RETBAR1, and RETBAR2, and are described below:

- RETCCW consisted of a 22-s blank period, 8 cycles of a 90° wedge rotating counter-clockwise with a period of 32 s, and a 22-s blank period. The duty cycle for a given point in the visual field was 25% (8 s of 32 s).
- RETCW was the same as RETCCW except that the wedge rotated clockwise.
- RETEXP consisted of a 22-s blank period, 8 cycles of a ring expanding away from the center of the screen with a period of 32 s, and a 22-s blank period. The last 4 s of each 32-s period was blank (thus helping distinguish foveal and peripheral responses). Ring size increased linearly with eccentricity. The duty cycle for a given point in the visual field was 19% (6 s of 32 s).
- RETCON was the same as RETEXP except that the ring contracted towards the center of the screen.
- RETBAR1 and RETBAR2 were identical, and consisted of a 16-s blank period, 4 bar movements lasting 32 s each (RIGHT, UP, LEFT, DOWN), a 12-s blank period, 4 bar movements lasting 32 s each (UPPER-RIGHT, UPPER-LEFT, LOWER-LEFT, LOWER-RIGHT), and a 16-s blank period. The capitalized term indicates the direction of bar movement. The last 4 s of each 32-s bar movement was blank (thus, the bar traversed the visual field in 28 s). The width of the bar was

1/8 of the full stimulus extent. The duty cycle for a given point in the visual field was 10% (3.11 s of 32 s).

Apertures were animated at a rate of 15 Hz, and each aperture was anti-aliased. On each aperture update, one of the 100 texture images was randomly selected (under the constraint that the same texture image is not presented consecutively) and presented within the confines of the aperture (using the continuous values of the aperture as opacity values). Each run consisted of 300 s x 15 Hz = 4500 stimulus frames.

Experimental design and task

A small semi-transparent dot (0.3° x 0.3°) at the center of the display was present throughout the experiment. The color of the central dot switched randomly to one of three colors (black, white, or red) every 1–5 s. Subjects were instructed to maintain fixation on the dot and to press a button whenever the color of the dot changed. The purpose of the task was to encourage fixation and allocation of attention to the center of the display. To further aid fixation, a semi-transparent fixation grid was superimposed on the display throughout the experiment (Schira et al., 2009).

Stimuli were presented using an NEC NP4000 projector. The projected image was focused onto a backprojection screen, and subjects viewed this screen via a mirror mounted on the RF coil. The projector operated at a resolution of 1024 x 768 @ 60 Hz. A Macintosh computer controlled stimulus presentation using code based on the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Behavioral responses were recorded using a Curdes FORP button box. Eyetracking was performed using an EyeLink 1000 system (SR Research, Mississauga, Ontario, Canada). Eyetracking data are available on ConnectomeDB for most subjects, but we caution that the quality of the data is variable due to obstructions within the head coil. Eyetracking data were not used in this paper.

The viewing distance to the backprojection screen was 101.5 cm, and the full stimulus extent (i.e. diameter of the circle within which apertures are shown) was 28.5 cm, yielding a total stimulus size of 16.0°. However, due to variations in subject setup, these numbers should be considered approximate. Furthermore, due to the confines of the MRI environment, some subjects were unable to see the very top and very bottom of the stimuli (approximately 1° at each end). This should be taken into account when interpreting the fMRI results.

pRF analysis

Population receptive field (pRF) model

We analyzed the time-series data of each grayordinate using a pRF model called the Compressive Spatial Summation model (Kay et al., 2013). This model is implemented in a toolbox called analyzePRF (http://cvnlab.net/analyzePRF/); to analyze the HCP 7T Retinotopy Dataset, we modified the implementation and archived the resulting code on the Open Science Framework web site (https://osf.io/bw9ec/).

The model predicts the fMRI time series as the sum of a stimulus-related time series and a baseline time series. The stimulus-related time series is obtained by computing the dot product between the stimulus apertures and a 2D isotropic Gaussian, applying a static power-law nonlinearity, scaling the result by a gain factor, and then convolving the result with a canonical hemodynamic response function (HRF). This can be expressed formally as $r(t) = (g \times (S(t) \bullet G)^n) * h(t)$ where $r(t)$ is the predicted stimulus-related time

series, g is a gain parameter, $S(t)$ is the stimulus aperture at time t, G is the 2D isotropic Gaussian, n is an exponent parameter, and $h(t)$ is a canonical HRF. This time series characterizes BOLD modulations driven by the stimulus. The baseline time series is obtained by computing a weighted sum of low-order polynomial terms (constant, linear, quadratic, etc.). This time series characterizes the baseline BOLD signal level, i.e., the MR signal intensity that is present in the absence of the stimulus.

The model yields several parameters of interest: two parameters (x, y) that indicate the position of the Gaussian, a parameter (σ) that indicates the standard deviation of the Gaussian, a parameter (n) that indicates the exponent of the power-law nonlinearity, and a parameter (g) that indicates the overall gain of the predicted responses. In pilot analyses, we found that the experimental paradigm used here generally does not provide enough statistical power to estimate the exponent parameter reliably. Thus, we did not attempt to estimate this parameter but instead fixed the value of n to 0.05, which is representative of typical values that we observed.

Stimulus pre-processing

Prior to model fitting, we performed pre-processing of the stimulus apertures. The original resolution of the apertures in each run is 768 pixels x 768 pixels x 4500 frames. Aperture values range between 0 and 1 where 0 indicates the absence of the texture image and 1 indicates the presence of the texture image. To reduce computational burden, we resized the apertures to 200 pixels x 200 pixels. Then, to match the temporal resolution of the stimulus to the temporal resolution of the fMRI data, we averaged consecutive groups of 15 frames. This yielded a final stimulus resolution of 200 pixels x 200 pixels x 300 frames. Model fitting was performed in pixel units, and model parameters were posthoc converted from pixel units to degrees by multiplying by a scaling factor of 16.0 deg / 200 pixels.

Model fitting

In pilot analyses of the fMRI data, we noticed the high propensity for local minima in model solutions. To reduce inaccuracies and biases due to local minima, we designed the following fitting approach. We first performed a grid fit in which a range of parameter combinations were evaluated. We densely sampled parameter space using 25 nonlinearly spaced eccentricity values between 0 and 16 degrees (0, 0.04, 0.09, 0.16, 0.22, 0.33, 0.43, 0.58, 0.73, 0.95, 1.2, 1.5, 1.8, 2.2, 2.7, 3.3, 3.9, 4.7, 5.6, 6.8, 8.0, 9.6, 11.3, 13.7, and 16 degrees), 32 angle values between 0 and 360 degrees (0, 11.25, 22.5, ..., and 348.75 degrees), and 13 size values on a log scale between 1 and 128 pixels (equivalent to 0.08, 0.11, 0.16, 0.23, 0.32, 0.45, 0.64, 0.91, 1.3, 1.8, 2.6, 3.6, 5.1, 7.2, and 10.2 degrees), yielding 25 x 32 x 13 = 10,400 parameter combinations. The combination yielding the closest fit (in a least-squares sense) to the data was identified. We then used this parameter combination as the initial seed in a nonlinear optimization procedure (MATLAB Optimization Toolbox, Levenberg-Marquardt algorithm). The gain parameter of the initial seed was set to 0.75 of the optimal-fitting parameter value to allow room for adjustment in the optimization. Also, the gain parameter was restricted to be non-negative to constrain the space of fits to solutions that predict positive BOLD responses to stimulation. Note that no spatial constraint (e.g. smoothing) was incorporated into the model fitting process. Thus, parameter estimates for grayordinates are independent, thereby maximizing resolution and minimizing bias.

We fit the pRF model not only to the data from each subject, but also to the data from three group-average subjects, which were constructed by averaging time-series data across subjects. One group-average subject is the result of averaging all 181 subjects (Subject 184); the second group-average subject is the result of averaging a randomly chosen half of the subjects (Subject 182); and the third group-average subject is the result of averaging the other half of the subjects (Subject 183). For each

individual subject and each group-average subject, we performed three separate model fits: one fit uses all six runs, a second fit uses only the first half of each of the six runs, and the third fit uses only the second half of each of the six runs. The rationale for these fits is that the first fit provides the best estimate of model parameters, whereas the second and third fits can be used to assess the reliability of model parameters.

Each fit produces six quantities of interest: angle, eccentricity, pRF size (calculated as $\sigma/\sqrt{n}$, which corresponds to one standard deviation of the Gaussian that describes responses to point images across the visual field (Kay et al., 2013)), gain, percentage of variance explained, and mean signal intensity (calculated as the mean of all time points). The dimensions of the final results are 184 subjects x 91,282 grayordinates x 3 model fits x 6 quantities (see Figure 2).

Surface visualization

The pre-processed time-series data (CIFTI format) reflect MSMAll-alignment of individual subjects to the fs_LR surface (Glasser et al., 2013; Robinson et al., 2018). The pRF model solutions are obtained by fitting each CIFTI grayordinate independently; thus, there are no additional spatial transformations applied. In this paper, we visualize pRF model solutions by mapping from fs_LR space to *fsaverage* using nearest-neighbor interpolation and then using orthographic projection to visualize the *fsaverage* surface (see Figure 3). The underlay for the group-average results is the thresholded *fsaverage* curvature map, whereas the underlay for individual-subject results is the curvature obtained from individual subjects. An alternative to the *fsaverage* curvature is to compute the average curvature of the 181 subjects, and this may be useful for certain investigations (see Supplementary Information).

Group-average maps in Figures 3–6 are thresholded at 9.8% variance explained. This threshold was determined by fitting a Gaussian Mixture Model with two Gaussians to the distribution of variance explained values across grayordinates in the group-average data and then identifying the value at which the posterior probability switches from the Gaussian with smaller mean to the Gaussian with larger mean. The interpretation of this procedure is that the Gaussian with smaller mean likely reflects noise (voxels that are not visually responsive), the Gaussian with larger mean likely reflects signal (voxels that are visually responsive), and values above the threshold are more likely to reflect signal than noise. The same procedure was performed for the distribution of variance explained values across grayordinates in individual-subject data, and this yielded a threshold of 2.2% variance explained. We used this more liberal threshold for the individual-subject maps in Figure 7.

Timing and behavioral analysis

Stimulus timing, scanner timing, and button presses were logged in a behavioral file for each run. Behavioral files are missing for a small fraction of the runs (17 of 1086) and button presses were not detected for one of the subjects (HCP subject ID 782561). Analysis of stimulus timing indicates that run durations were highly reliable: the central 95% of run durations lie within the range [299.974 s, 299.982 s]. Analysis of scanner timing indicates that synchronization of the stimulus computer and the scanner was robust, with the exception of two runs in which scanner acquisition may have started 1 s (1 TR) too early relative to the stimulus.

To quantify behavioral performance, we calculated for each run, (A–B)/C x 100, where A indicates the number of successful detections of color changes (defined as the existence of a button press within 1 s of a color change), B indicates the number of extraneous button presses, and C indicates the total number

of color changes. We then averaged this performance value across the six runs. Behavioral performance was quite good overall: the central 50% of performance values across subjects lie within the range [91.6%, 97.8%].

References

Abdollahi RO, Kolster H, Glasser MF, Robinson EC, Coalson TS, Dierker D, Jenkinson M, Van Essen DC, Orban GA (2014) Correspondences between retinotopic areas and myelin maps in human visual cortex. NeuroImage 99:509–524.

Amunts K, Malikovic A, Mohlberg H, Schormann T, Zilles K (2000) Brodmann's areas 17 and 18 brought into stereotaxic space-where and how variable? NeuroImage 11:66–84.

Arcaro MJ, McMains SA, Singer BD, Kastner S (2009) Retinotopic organization of human ventral visual cortex. J Neurosci 29:10638–10652.

Arcaro MJ, Pinsk MA, Kastner S (2015) The Anatomical and Functional Organization of the Human Visual Pulvinar. J Neurosci 35:9848–9871.

Barch DM et al. (2013) Function in the human connectome: task-fMRI and individual differences in behavior. NeuroImage 80:169–189.

Benson NC, Butt OH, Brainard DH, Aguirre GK (2014) Correction of distortion in flattened representations of the cortical surface allows prediction of V1-V3 functional organization from anatomy. Einhäuser W, ed. PLoS computational biology 10:e1003538.

Benson NC, Butt OH, Datta R, Radoeva PD, Brainard DH, Aguirre GK (2012) The retinotopic organization of striate cortex is well predicted by surface topology. Curr Biol 22:2081–2085.

Bock AS, Binda P, Benson NC, Bridge H, Watkins KE, Fine I (2015) Resting-State Retinotopic Organization in the Absence of Retinal Input and Visual Experience. J Neurosci 35:12366–12382.

Brainard DH (1997) The Psychophysics Toolbox. Spat Vis 10:433–436.

Cotton PL, Smith AT (2007) Contralateral visual hemifield representations in the human pulvinar nucleus. Journal of neurophysiology 98:1600–1609.

DeSimone K, Viviano JD, Schneider KA (2015) Population Receptive Field Estimation Reveals New Retinotopic Maps in Human Subcortex. J Neurosci 35:9836–9847.

DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, Miller D, Neitz J (1996) Mapping striate and extrastriate visual areas in human cerebral cortex. Proceedings of the National Academy of Sciences of the United States of America 93:2382–2386.

Dougherty RF, Koch VM, Brewer AA, Fischer B, Modersitzki J, Wandell B (2003) Visual field representations and locations of visual areas V1/2/3 in human visual cortex. Journal of vision 3:586–598.

Dumoulin SO, Wandell B (2008) Population receptive field estimates in human visual cortex. NeuroImage 39:647–660.

Engel SA, Glover GH, Wandell B (1997) Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cereb Cortex 7:181–192.

Engel SA, Rumelhart DE, Wandell B, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. Nature 369:525.

Epstein R, Kanwisher N (1998) A cortical representation of the local visual environment. Nature 392:598–

601.

Glasser MF, Coalson TS, Bijsterbosch JD, Harrison SJ (2017) Using Temporal ICA to Selectively Remove Global Noise While Preserving Global Signal in Functional MRI Data. bioRxiv.

Glasser MF, Coalson TS, Robinson EC, Hacker CD, Harwell J, Yacoub E, Ugurbil K, Andersson J, Beckmann CF, Jenkinson M, Smith SM, Van Essen DC (2016) A multi-modal parcellation of human cerebral cortex. Nature 536:171–178.

Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR, Van Essen DC, Jenkinson M, WU-Minn HCP Consortium (2013) The minimal preprocessing pipelines for the Human Connectome Project. NeuroImage 80:105–124.

Glasser MF, Van Essen DC (2011) Mapping human cortical areas in vivo based on myelin content as revealed by T1- and T2-weighted MRI. J Neurosci 31:11597–11616.

Griffanti L, Salimi-Khorshidi G, Beckmann CF, Auerbach EJ, Douaud G, Sexton CE, Zsoldos E, Ebmeier KP, Filippini N, Mackay CE, Moeller S, Xu J, Yacoub E, Baselli G, Ugurbil K, Miller KL, Smith SM (2014) ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. NeuroImage 95:232–247.

Grill-Spector K, Weiner KS (2014) The functional architecture of the ventral temporal cortex and its role in categorization. Nature Rev Neurosci 15:536–548.

Hinds OP, Rajendran N, Polimeni JR, Augustinack JC, Wiggins G, Wald LL, Diana Rosas H, Potthast A, Schwartz EL, Fischl B (2008) Accurate prediction of V1 location from cortical folds in a surface coordinate system. NeuroImage 39:1585–1599.

Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. J Neurosci 17:4302–4311.

Kastner S, De Weerd P, Pinsk MA, Elizondo MI, Desimone R, Ungerleider LG (2001) Modulation of sensory suppression: implications for receptive field sizes in the human visual cortex. Journal of neurophysiology 86:1398–1411.

Kastner S, DeSimone K, Konen CS, Szczepanski SM, Weiner KS, Schneider KA (2007) Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. Journal of neurophysiology 97:3494–3507.

Katyal S, Zughni S, Greene C, Ress D (2010) Topography of covert visual attention in human superior colliculus. Journal of neurophysiology 104:3074–3083.

Kay KN, Winawer J, Mezer A, Wandell B (2013) Compressive spatial summation in human visual cortex. Journal of neurophysiology 110:481–494.

Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, Tanaka K, Bandettini PA (2008) Matching categorical object representations in inferior temporal cortex of man and monkey. Neuron 60:1126–1141.

Larsson J, Heeger DJ (2006) Two retinotopic visual areas in human lateral occipital cortex. J Neurosci 26:13128–13142.

Mackey WE, Winawer J, Curtis CE (2017) Visual field map clusters in human frontoparietal cortex. Elife 6:2704.

Malach R, Levy I, Hasson U (2002) The topography of high-order human object areas. Trends in cognitive sciences 6:176–184.

Maunsell JH, Van Essen DC (1983) The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. J Neurosci 3:2563–2586.

McKeefry DJ, Watson JD, Frackowiak RS, Fong K, Zeki S (1997) The activity in human areas V1/V2, V3, and V5 during the perception of coherent and incoherent motion. NeuroImage 5:1–12.

Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat Vis 10:437–442.

Press WA, Brewer AA, Dougherty RF, Wade AR, Wandell B (2001) Visual areas and spatial summation in human visual cortex. Vision research 41:1321–1332.

Raemaekers M, Schellekens W, van Wezel RJA, Petridou N, Kristo G, Ramsey NF (2014) Patterns of resting state connectivity in human primary visual cortical areas: a 7T fMRI study. NeuroImage 84:911–921.

Robinson EC et al. (2018) Multimodal surface matching with higher-order smoothness constraints. NeuroImage 167:453–465.

Robinson EC, Jbabdi S, Glasser MF, Andersson J, Burgess GC, Harms MP, Smith SM, Van Essen DC, Jenkinson M (2014) MSM: a new flexible framework for Multimodal Surface Matching. NeuroImage 100:414–426.

Salimi-Khorshidi G, Douaud G, Beckmann CF, Glasser MF, Griffanti L, Smith SM (2014) Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. NeuroImage 90:449–468.

Schira MM, Tyler CW, Breakspear M, Spehar B (2009) The foveal confluence in human visual cortex. J Neurosci 29:9050–9058.

Schneider KA, Kastner S (2005) Visual responses of the human superior colliculus: a high-resolution functional magnetic resonance imaging study. Journal of neurophysiology 94:2491–2503.

Schneider KA, Richter MC, Kastner S (2004) Retinotopic organization and functional subdivisions of the human lateral geniculate nucleus: a high-resolution functional magnetic resonance imaging study. J Neurosci 24:8975–8985.

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.

Silver MA, Kastner S (2009) Topographic maps in human frontal and parietal cortex. Trends in cognitive sciences 13:488–495.

Silver MA, Ress D, Heeger DJ (2005) Topographic maps of visual spatial attention in human parietal cortex. Journal of neurophysiology 94:1358–1371.

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.

Swisher JD, Halko MA, Merabet LB, McMains SA, Somers DC (2007) Visual topography of human intraparietal sulcus. J Neurosci 27:5326–5337.

Tootell RB, Dale AM, Sereno MI, Malach R (1996) New images from human visual cortex. Trends in neurosciences 19:481–489.

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.

Ugurbil K et al. (2013) Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. NeuroImage 80:80–104.

Van Essen DC et al. (2012) The Human Connectome Project: a data acquisition perspective. NeuroImage 62:2222–2231.

Van Essen DC, Smith J, Glasser MF, Elam J, Donahue CJ, Dierker DL, Reid EK, Coalson T, Harwell J (2017) The Brain Analysis Library of Spatial maps and Atlases (BALSA) database. NeuroImage 144:270–274.

Van Essen DC, Smith SM, Barch DM, Behrens TEJ, Yacoub E, Ugurbil K, WU-Minn HCP Consortium (2013) The WU-Minn Human Connectome Project: an overview. NeuroImage 80:62–79.

Vu AT, Jamison K, Glasser MF, Smith SM, Coalson T, Moeller S, Auerbach EJ, Ugurbil K, Yacoub E (2016) Tradeoffs in pushing the spatial resolution of fMRI for the 7T Human Connectome Project. NeuroImage.

Wade AR, Brewer AA, Rieger JW, Wandell B (2002) Functional measurements of human ventral occipital cortex: retinotopy and colour. Philosophical transactions of the Royal Society of London 357:963–973.

Wandell B, Brewer AA, Dougherty RF (2005) Visual field map clusters in human cortex. Philosophical transactions of the Royal Society of London 360:693–707.

Wandell B, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. Neuron 56:366–383.

Wandell B, Winawer J (2011) Imaging retinotopic maps in the human brain. Vision research 51:718–737.

Wandell B, Winawer J (2015) Computational neuroimaging and population receptive fields. Trends in cognitive sciences 19:349–357.

Wang L, Mruczek REB, Arcaro MJ, Kastner S (2015) Probabilistic Maps of Visual Topography in Human Cortex. Cereb Cortex 25:3911–3931.

Supplementary Information

Supplementary Movies 1–12: pRF results on dynamic rotating cortical surfaces

Each movie shows group-average pRF model solutions on the inflated *fsaverage* surface. Format same as in Figure 4. There are a total of 2 hemispheres x 6 maps (curvature, angle, eccentricity, zoomed eccentricity, pRF size, variance explained) = 12 movies. The movies are accessible at the OSF web site (https://osf.io/bw9ec/).

Supplementary Discussion 1: The HCP 7T Retinotopy Dataset in BALSA and Connectome Workbench

1. Introduction.

The HCP 7T retinotopy dataset is freely available via the BALSA database as downloadable scene files that can be viewed using Connectome Workbench. Each scene file contains multiple individual scenes that display retinotopic and associated data in an arrangement that provides a useful starting point for visualization and further analysis. After downloading, scenes can be displayed in 'wb_view', but **you will need Connectome Workbench v. 1.3 or higher** (https://www.humanconnectome.org/software/get-connectome-workbench), **because the scenes rely on several recently incorporated features.**

The BALSA URL for this study is https://balsa.wustl.edu/study/show/9Zkk.
The primary scene file ("Retinotopy_HCP_7T_181_Fit1.scene"; https://balsa.wustl.edu/sceneFile/show/kj48) shows retinotopy results for all 181 7T subjects with retinotopy data displayed on various geometries (spheres, flatmaps, etc.) on the "fs_LR" surface mesh aligned using the "MSMAll" areal feature-based alignment. We encourage investigators to use this dataset, which is in the original space in which the retinotopic analysis was carried out and also the original space of the HCP_MMP1.0 multimodal cortical parcellation and a growing number of other datasets. The primary scene file shows the retinotopic analysis fit for the whole dataset (fit1). Additional scene files provide results from (i) analyzing two subgroups of retinotopy subjects (90- and 91-subject groups), (ii) Data from the first half of each run of each subject (fit2) and the second half of each run of each subject (fit3), and (iii) retinotopy results mapped to the fsaverage atlas space.

We briefly describe the layout of the first scene in the primary scene file and give useful navigation tips, as this scene has commonalities with other scenes in this and the other scene files. We then show figure previews for the remaining scenes in the primary scene file, followed by a listing of additional scene files for the 90 and 91 subject datasets and the fit2 and fit3 analyses, plus illustrations of the two scenes in the fsaverage scene file illustrating some of the limitations of mapping the data to fsaverage using nearest neighbor interpolation and displaying it on the fsaverage folding pattern as is done in the main text.

Retinotopy_HCP_7T_181_Fit1.scene: https://balsa.wustl.edu/sceneFile/show/kj48

2. Viewing and navigating the first retinotopy scene.

Figure S1 shows the group average (n = 181) maps of all of the retinotopic data (fit1) polar angle, eccentricity, pRF size, and myelin viewed on spherical maps of the left and right hemisphere, oriented for maximal coverage of retinotopic areas (slightly different than the orientation used for the main text figures). The HCP_MMP1.0 borders appear as an overlay (white contours); the underlay is the group average folding map (FreeSurfer's curvature measure) from the same 181 subjects aligned with the same MSMAll registration algorithm; this underlay is visible only in transparent regions where the model R^2 is less than 9.8%. Using the average folding map derived from the same subjects used for the retinotopy analysis provides a more realistic representation of folding-function relationships in this dataset than if one uses the fsaverage map of average folding from a separate set of subjects aligned using a different registration algorithm (see Figure S10 for a direct comparison).

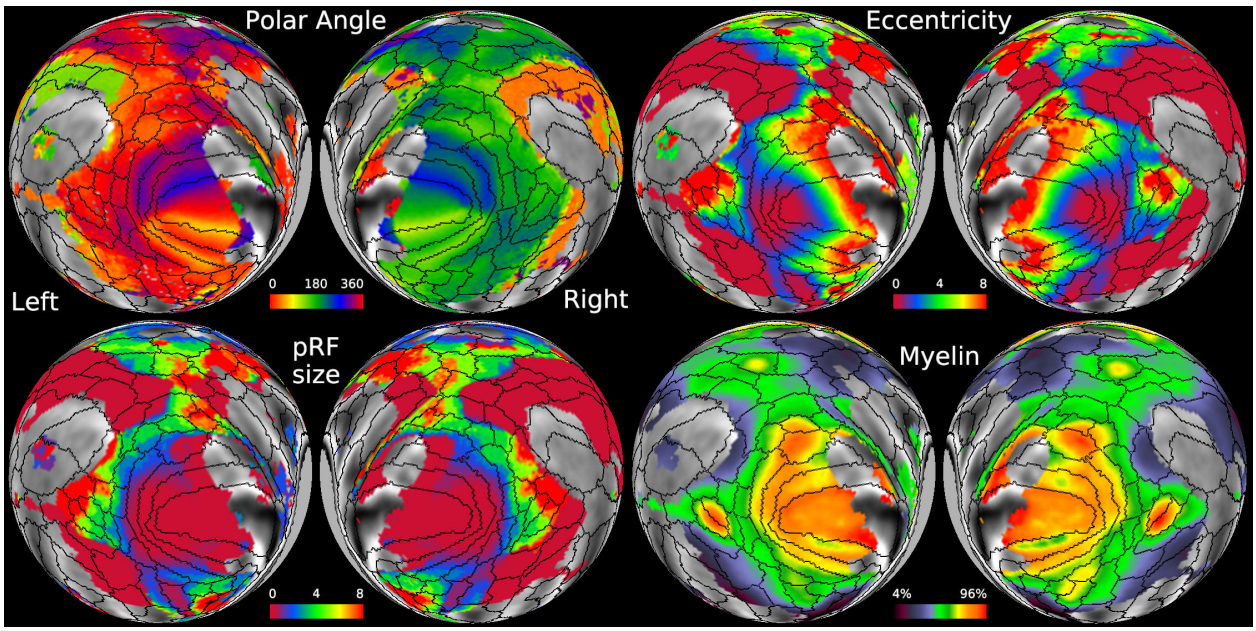

Figure S1. Group average polar angle, eccentricity, pRF size, and myelin maps viewed on spherical maps with the light shading omitted in order to make the view look flatter. For the polar angle map (0 to 360), 0 (and 360) degrees corresponds to the horizontal meridian in the right visual hemifield (left hemisphere), whereas 180 degrees corresponds to the horizontal meridian in the left visual hemifield. Thresholding of all maps is at 9.8% of the variance explained for the group in this and other group average figures, showing the group average folding map underneath the thresholded portions. Data at https://balsa.wustl.edu/k0Z2.

Switching from Tile Tabs to single tab views. This scene opens in "Tile Tabs" mode, showing four tabs concurrently. To view in single-tab mode, select the tab of interest then select View: Exit Tile Tabs or use a keystroke shortcut – CMD-M (Mac) or CTRL-M (Linux). To return to Tile

Tabs, select View: Enter Tile Tabs, repeat the CMD-M/CTRL-M shortcut, or press “Show” in the Scenes dialog window.

View other maps using the Overlay Toolbox. Each scene includes many loaded but initially invisible files that can be displayed using the Overlay Toolbox. For example, instead of the default polar angle “0 – 360” palette (0 = red = right hemifield horizontal meridian, with distinct hues for the left and right visual hemifield), a mirror-symmetric palette (analogous to Fig. 4 of the main text) can be viewed by toggling off the top polar angle layer in the Overlay Toolbox menu. Toggling off the third layer (“PolarAngleMirror”) reveals the complete average curvature map in layer 4. For each layer, the File pulldown menu enables selection of any of the currently loaded maps (mostly in ‘dscalar’ format).

Comparisons with the Wang2015 retinotopic parcellation. The Features Toolbox (on the right, open by default for this scene) enables turning off of the HCP_MMP1.0 areal borders (uncheck the two checked boxes labeled “Q1-Q6_RelatedParcellation210.R…” and “Q1-Q6_RelatedParcellation210.L…”) and turning on the Wang2015 retinotopic borders (“wang2015.L…..” and “wang2015.R…”).

3. Additional scenes in the primary scene file.

Figures S2 – S4 show retinotopy data on flatmaps, inflated surfaces, and average midthickness surfaces in the same layout as for Figure S1.

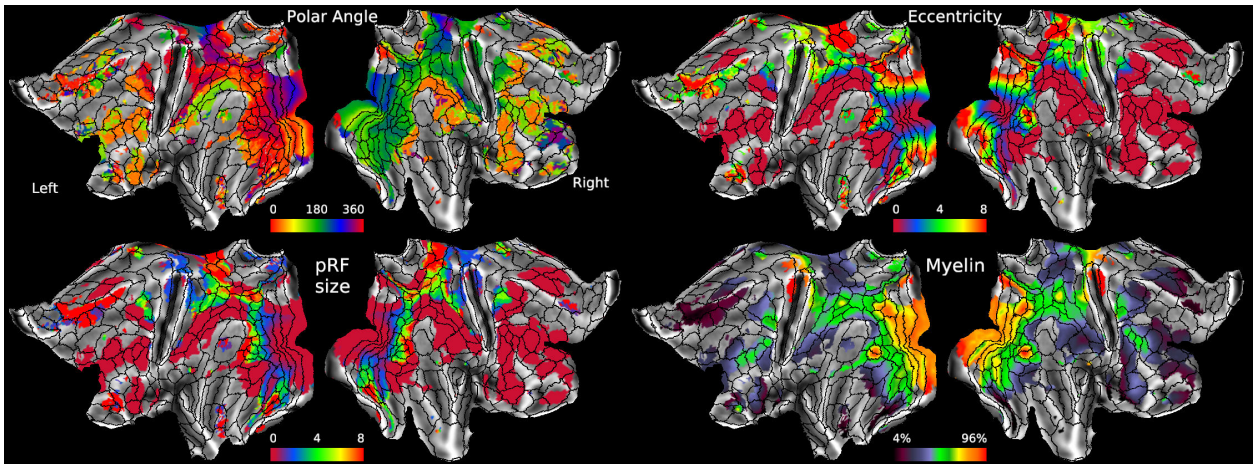

Figure S2. Group average polar angle, eccentricity, pRF size, and myelin maps viewed on flat maps of the left and right hemispheres. Data at https://balsa.wustl.edu/0gVl.

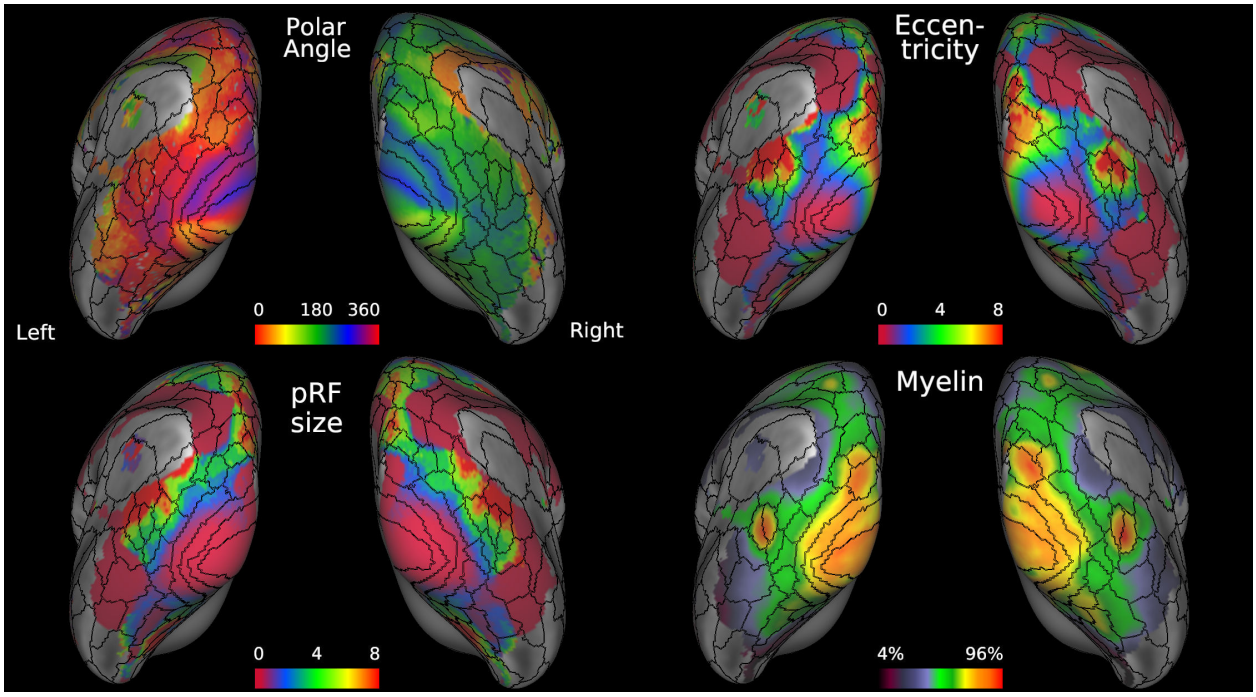

Figure S3. Group average polar angle, eccentricity, pRF size, and myelin maps viewed on inflated left and right hemispheres. Data at https://balsa.wustl.edu/2BzG.

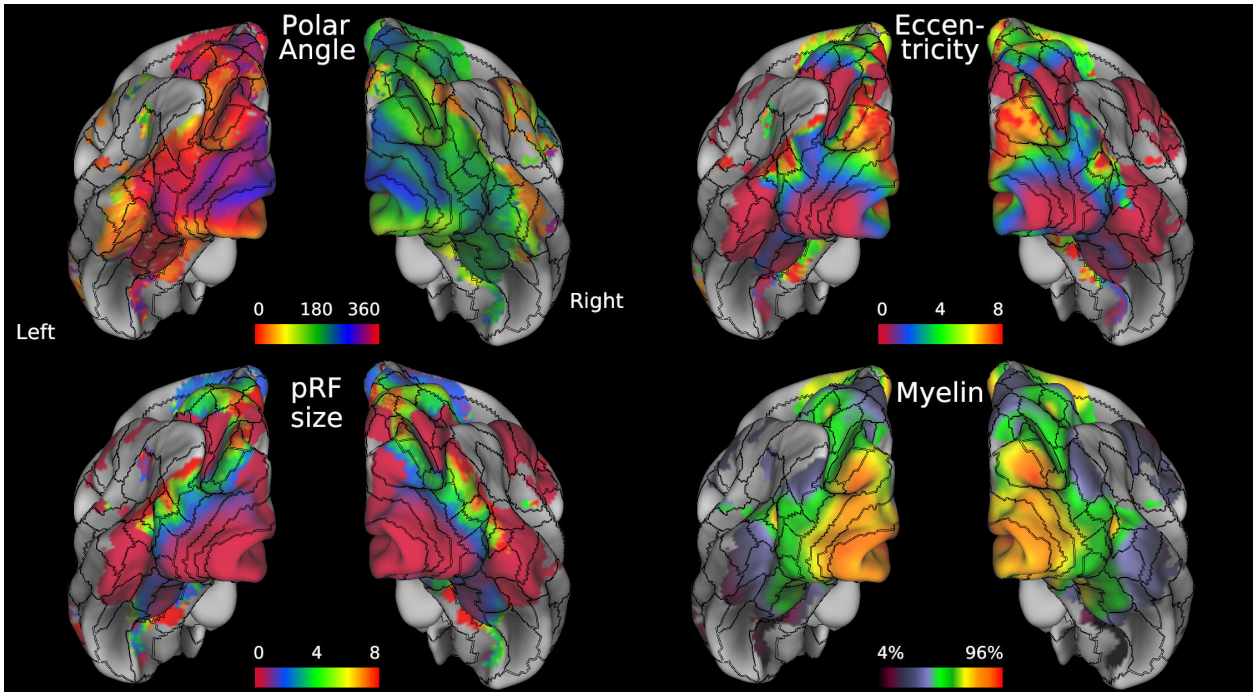

Figure S4. Group average polar angle, eccentricity, pRF size, and myelin maps viewed on average left and right midthickness surfaces. Data at https://balsa.wustl.edu/r15G.

Subcortical retinotopy maps. Figure S5 shows group average retinotopy maps on subcortical volume slices.

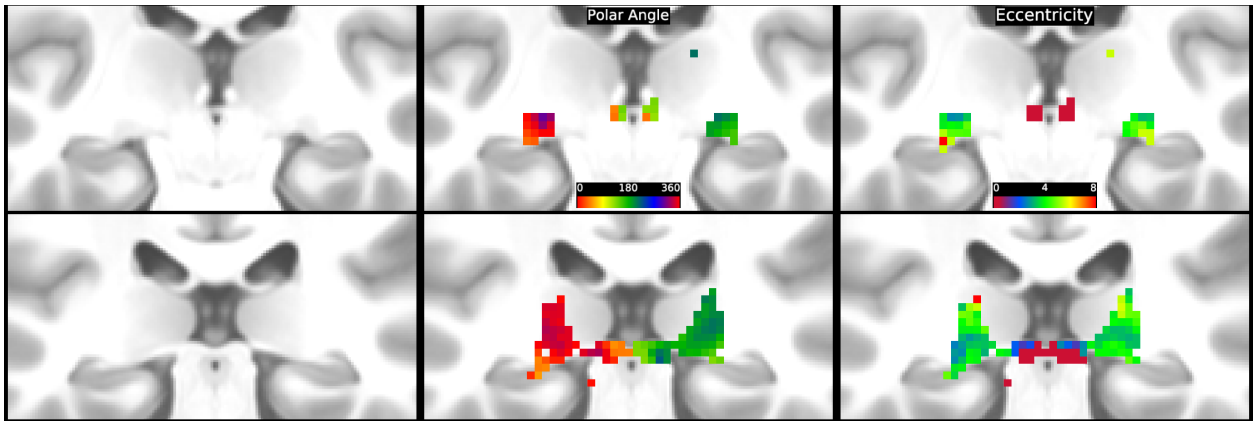

Figure S5. Group average polar angle (middle) and eccentricity (right) overlaid on average T1w coronal slices (left) at the level of the LGN (top row, y = -25) and at the level of the superior colliculus and pulvinar (bottom row, y = -30), as in Fig. 6 of main text. Data at https://balsa.wustl.edu/xlq2.

Individual-subject maps. Figure S6 shows retinotopic and myelin maps for HCP subject 100610 (lowest subject number of the 181 studied), displayed on spherical maps in the same arrangement as Figure S1 and akin to Figure 7 of the main text. Retinotopy for each individual can be examined and compared using the scroll bar to the left of the Map layer for PolarAngle; since the maps are yoked in this scene, changing one map changes those in the other tabs as well.

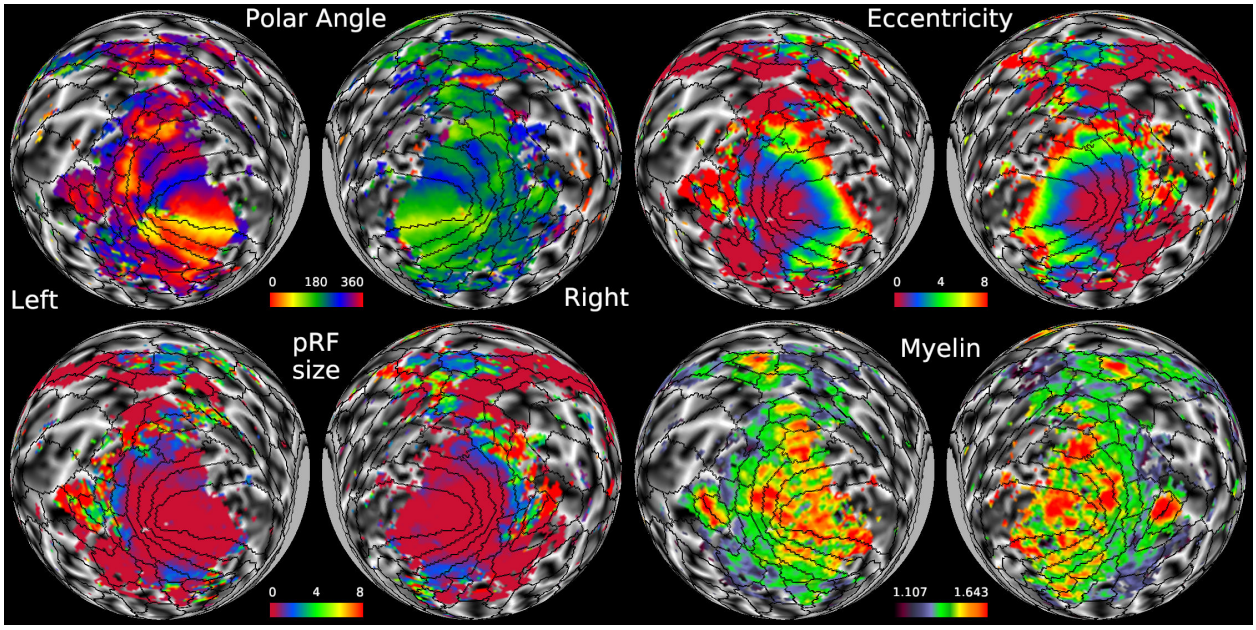

Figure S6. Individual-subject polar angle, eccentricity, pRF size, and myelin maps viewed on top of individual subject folding maps for the left and right spherical surfaces. Thresholding for all individual subjects is at 2.2% of the variance explained for that subject. https://balsa.wustl.edu/PL72.

A nonexhaustive list of subjects with reasonable signal quality but potentially atypical polar angle retinotopy (organization not just alignment to average) include: map 5 (subject 105923, blue = inf VM), 7 (109123, ditto), 9 (111514, ditto), 15 (125525, upper and lower), 41, 43 (LH – foveal green ipsi HM), 81, 86, 91, 105, 174.

An alternate arrangement for viewing retinotopy. Figure S7 shows group average retinotopy on spherical maps with right hemispheres all in the top row and left hemispheres in the bottom row. This can be convenient when inspecting data in the Tile Tabs mode, but it has a less convenient aspect ratio upon switching to single tab mode.

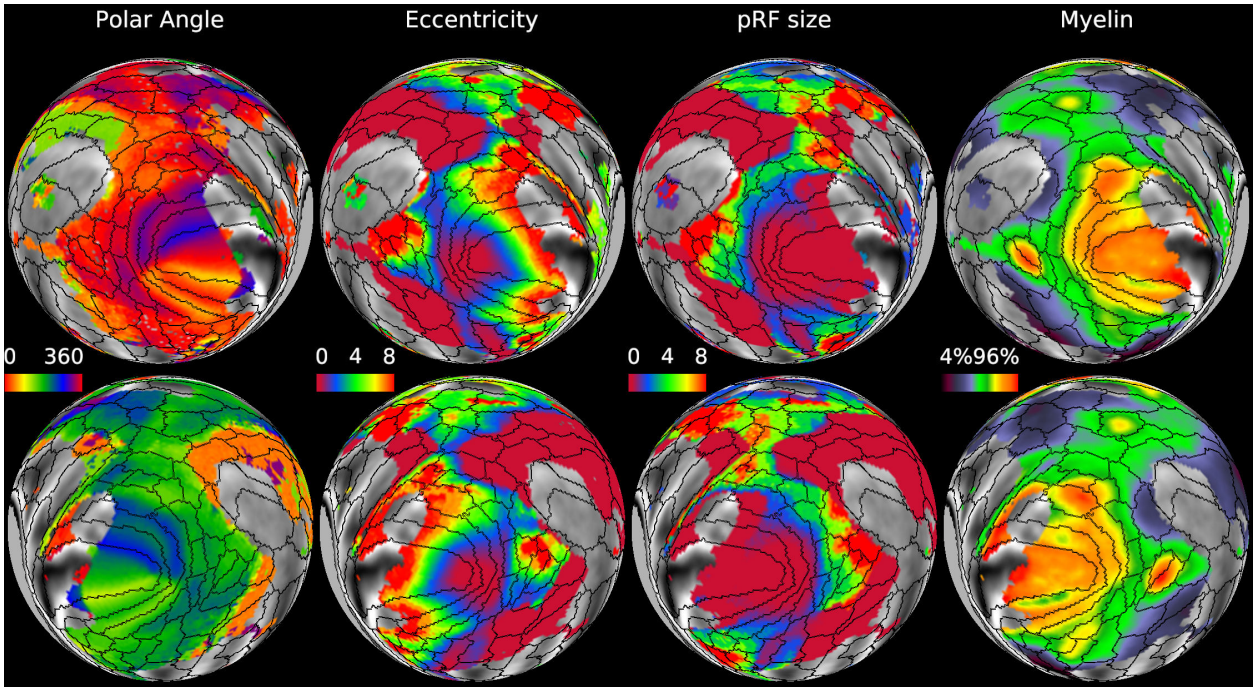

Figure S7. Group average polar angle, eccentricity, pRF size, and myelin maps viewed on average left and right spheres in a left on top; right on bottom configuration. Data at https://balsa.wustl.edu/7ZMj.

Comparing HCP_MMP1.0 and Wang et al. (2015) parcellations. Figure S8 shows the Wang et al. (2015) retinotopic parcellation (colored parcels) compared to the Glasser et al. (2016) HCP_MMP1.0 parcellation (black borders). There is good agreement for early visual areas (V1 – V3 and V3A) bilaterally, aside from the retinotopic parcel extent reflecting incomplete mapping of the visual field. For other areas the alignment is modest to poor and likely results from differences in alignment method and differences in parcellation (e.g. around V4, see Glasser et al. 2016).

Retinotopy_HCP_7T_Atlas.scene: https://balsa.wustl.edu/sceneFile/show/vrVl

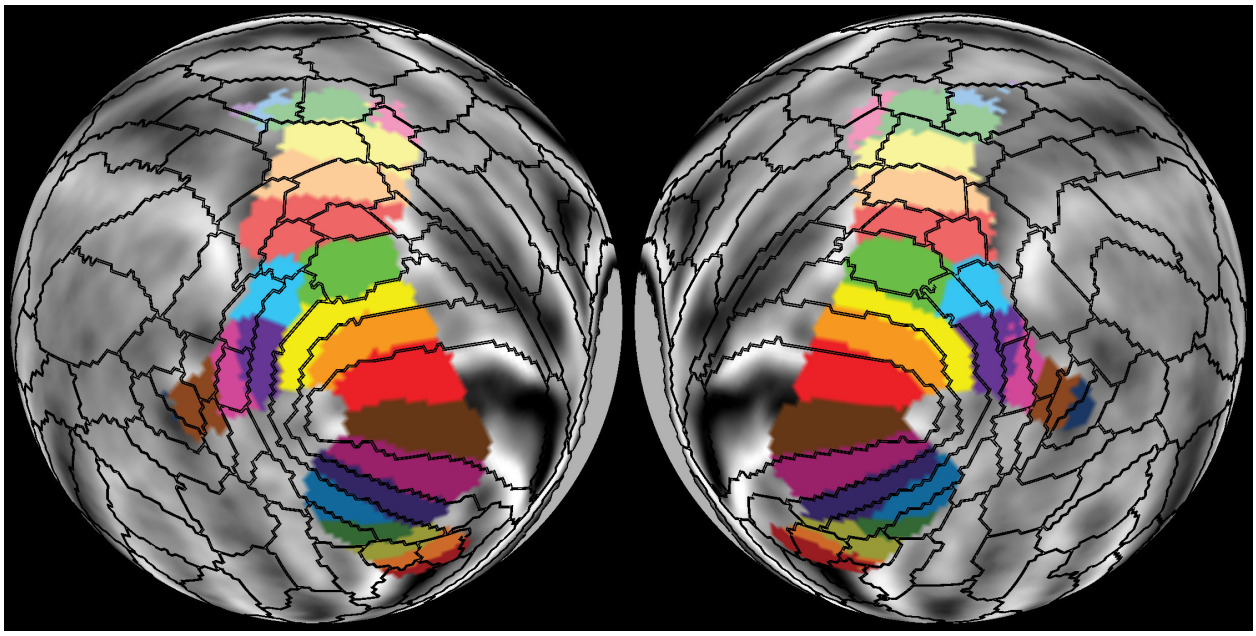

Figure S8. Wang et al. (2015) retinotopic parcellation (colored parcels) compared to Glasser et al. (2016) HCP_MMP1.0 parcellation (black borders). Data at https://balsa.wustl.edu/xlxx.

4. Fit2 and Fit3 retinotopic data.

The scene files for the fit2 (first half of each run of each subject) and fit3 (second half of each run of each subject) retinotopic analyses (see Methods in main text) replicate the scenes shown above (Figures S1 – S7). Here we provide BALSA URLs for the corresponding scene files; linking to these provides access to the individual scene URLs as well.

Retinotopy_HCP_7T_181_Fit2.scene: https://balsa.wustl.edu/sceneFile/show/m4Zp

Retinotopy_HCP_7T_181_Fit3.scene: https://balsa.wustl.edu/sceneFile/show/1zjl

5. 90-subject and 91-subject retinotopic data.

The scene files for the 90-subject and 91-subject subgroups of retinotopic analyses (see Methods in main text) replicate the scenes shown above (Figures S1 – S5 and S7). Here we provide BALSA URLs for the corresponding scene files; linking to these provides access to the individual scene URLs as well.

Retinotopy_HCP_7T_90_Fit1.scene: https://balsa.wustl.edu/sceneFile/show/6zz7

Retinotopy_HCP_7T_91_Fit1.scene: https://balsa.wustl.edu/sceneFile/show/N8Xj

6. Retinotopy and folding data mapped to fsaverage atlas.

Figures S9 and S10 show data mapped from the fs_LR 32k mesh to the hemisphere-specific fsaverage (fs_L and fs_R) atlas surfaces. They are in scene file:

Retinotopy_HCP_7T_181_fsaverage.scene: https://balsa.wustl.edu/sceneFile/show/281D

Mean curvature for fsaverage vs MSMAll. Figure S9 shows maps of mean curvature (folding) computed from the fsaverage atlas (column 1) and the average across the 181 7T retinotopy subjects aligned using MSMAll (column 2). Columns 3 and 4 show the same maps binarized for positive (gyral) vs negative (sulcal) average folding values. The lower values of mean curvature for the MSMAll reflect its greater reliance on using areal features for alignment, in contrast to fsaverage reliance on alignment based only on folding even in regions of high folding variability. Additionally, the fsaverage curvature map appears to have been recalculated on the fsaverage average surface itself, rather than being the average of the curvature maps across the subjects that generated fsaverage. Data are displayed on the fsaverage ("fs_L" and "fs_R") surface meshes. Although some features are grossly similar (e.g. calcarine sulcus), many are quite different between the two folding maps. Thus, it can be misleading to interpret the relationships between retinotopic data and folds of different datasets aligned using different registration algorithms or folding maps computed from an average surface vs the average map across subjects.

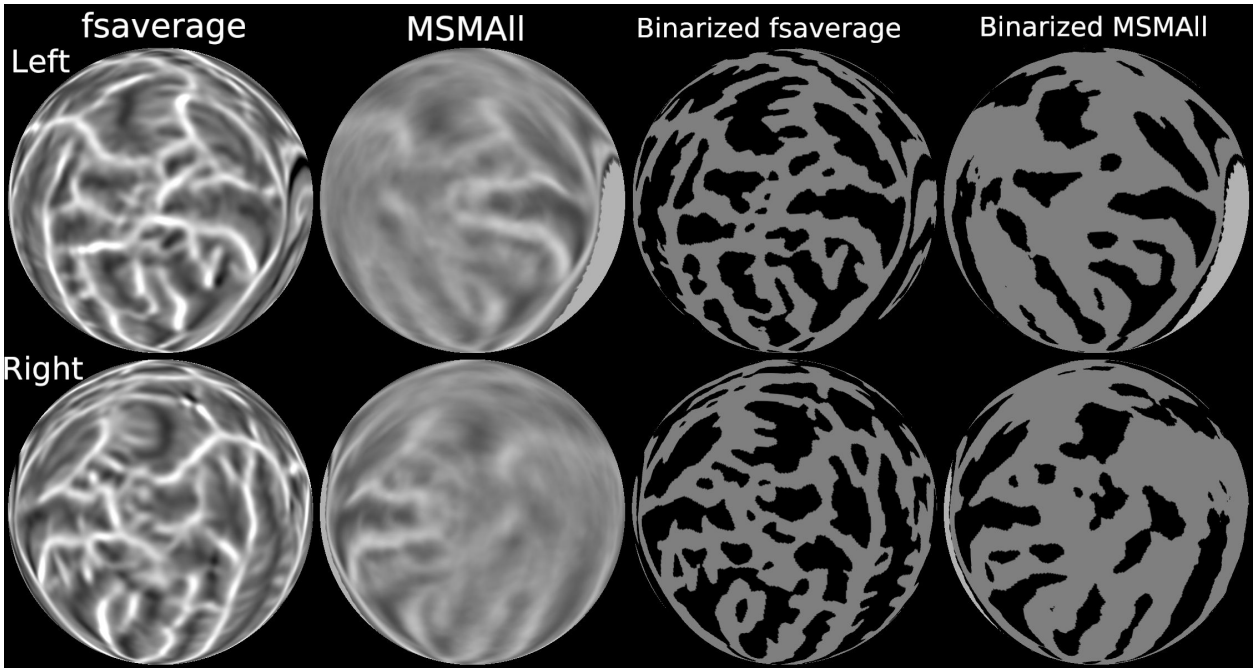

Figure S9. Maps of mean curvature (folding) computed from fsaverage (column 1) and from the 181 7T retinotopy subjects aligned using MSMAll (column 2).

Columns 3 and 4 show the same maps binarized for positive (gyral) vs negative (sulcal) average folding values. Data are displayed on the fsaverage ("fs_L" and "fs_R") surface meshes. Top row: left hemisphere spheres. Bottom row: right hemisphere spheres. Data at https://balsa.wustl.edu/VN4g.

Figure S10 shows the retinotopy data from the fs_LR 32k mesh to the hemisphere-specific fsaverage (fs_L and fs_R) atlas surfaces using two methods. One was the nearest neighbor method as described in the main Methods, which suffers from "pixelation" when upsampled from the fs_LR 32k mesh on which the data were computed to the 164k fs_L and fs_R fsaverage meshes. The other was an adaptive barycentric method, which is inherently smoother. However, for the polar angle and eccentricity maps, we first converted to Euclidean coordinates, in order to avoid discontinuities arising solely from representation, followed by adaptive barycentric area resampling, then conversion back into polar angle and eccentricity.

Note that the fsaverage datasets are much larger (and slower to load) because they were upsampled to the 164k FreeSurfer mesh, rather than the fs_LR 32k mesh in which the retinotopy analyses were carried out.

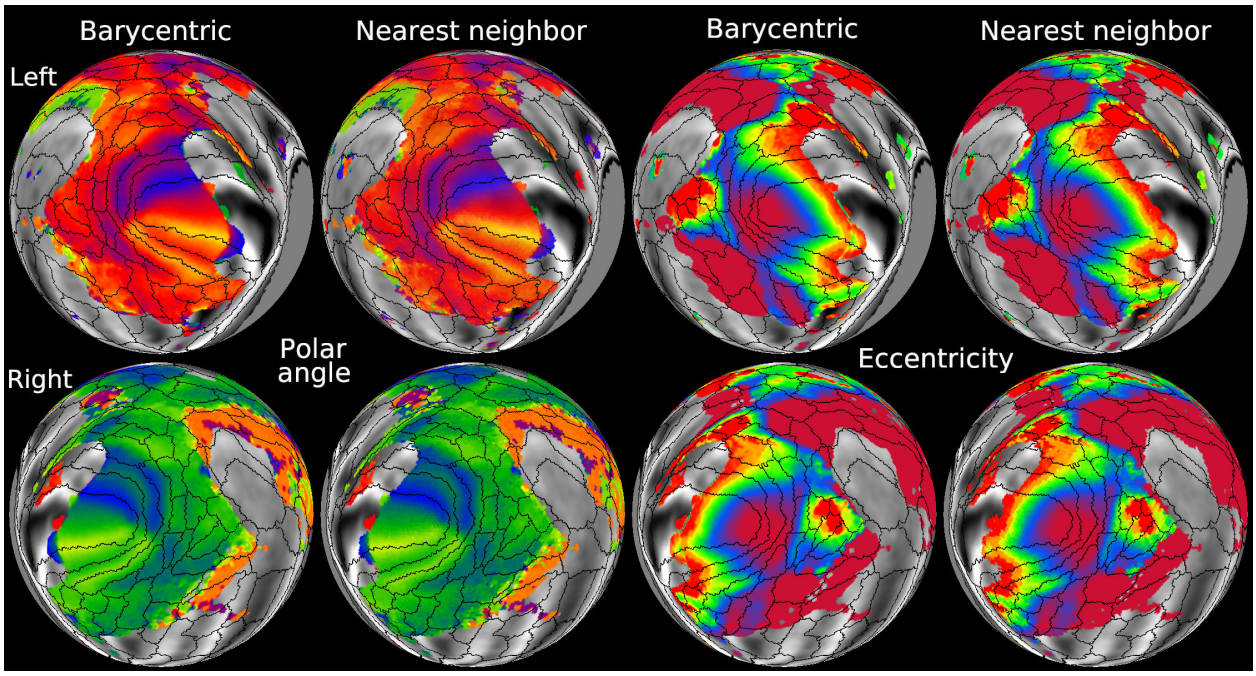

Figure S10. Group average polar angle (columns 1 and 2) and eccentricity (columns 3 and 4) maps viewed on fsaverage spheres. Thresholding is at 9.8% of the variance explained. Columns 1 and 3: Barycentric mapping from HCP fs_LR to fsaverage mesh. Columns 2 and 4: Nearest neighbor mapping from HCP fs_LR to fsaverage mesh. HCP_MMP1.0 areal borders resampled to fsaverage are displayed in all panels. Data at https://balsa.wustl.edu/3zNv.