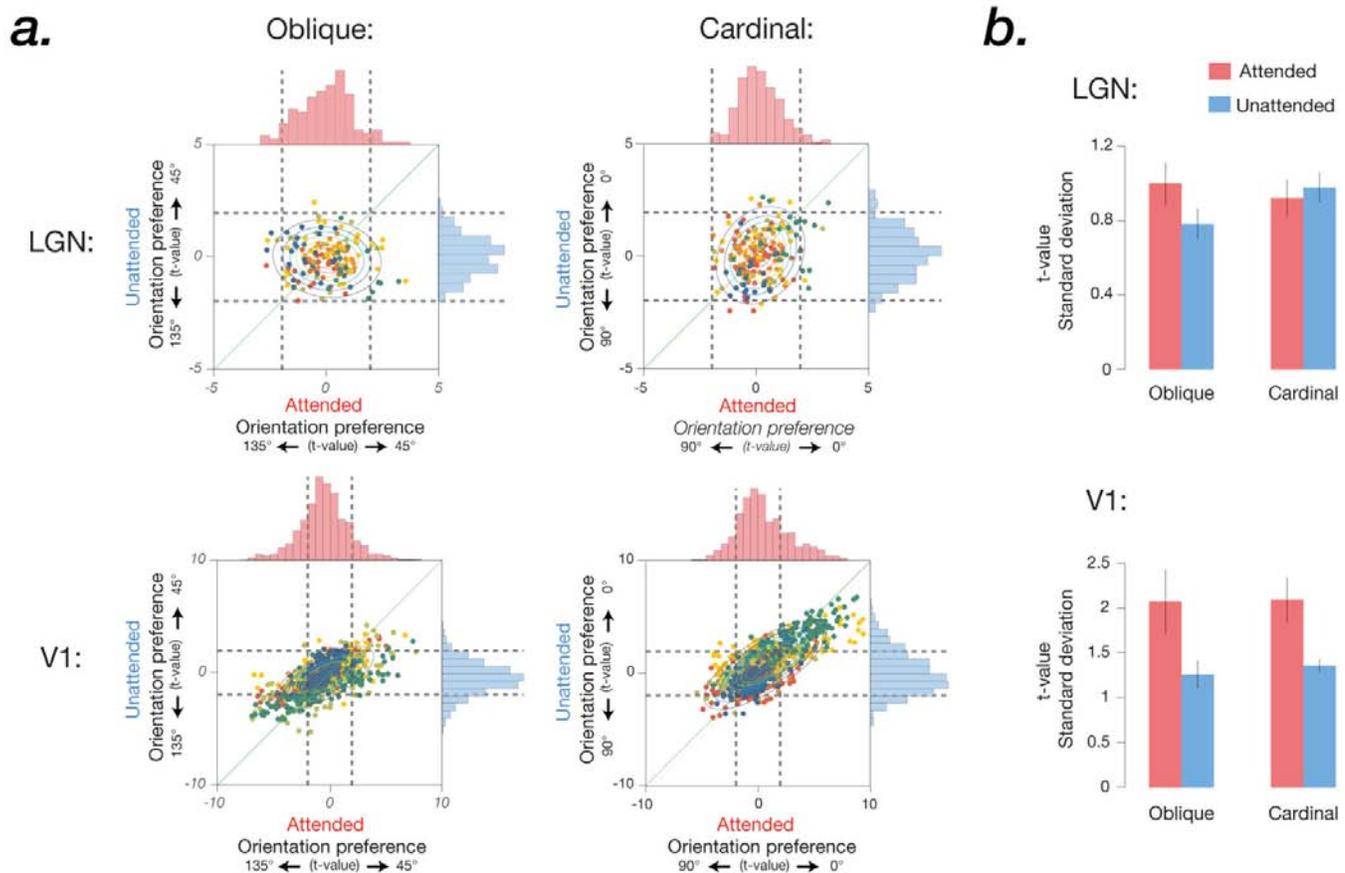


Supplementary Figure 1

**LGN ROI selection**

a) Example coronal (left) and sagittal (right) slices from proton density-weighted images for two subjects, and the corresponding anatomically defined LGN ROI (in blue). Images obtained using this sequence reveal a teardrop-shaped structure, often with a gap corresponding to the different density of the white matter sitting between the pulvinar and the LGN, which can be seen in the sagittal slices. The selection of voxels from this anatomical localizer was conservatively lateralized, to avoid including the adjacent lateral inferior pulvinar. b) Decoding results and BOLD response when strictly considering the functional localizer for ROI selection. The patterns of results are qualitatively quite similar to those using our anatomically restricted ROI (Fig. 1B, C). A repeated measures ANOVA on the LGN decoding data revealed a significant interaction between attention x orientation x area ( $F(1,47)=6.71, p<.05$ ), much like our PD/functional localizer.

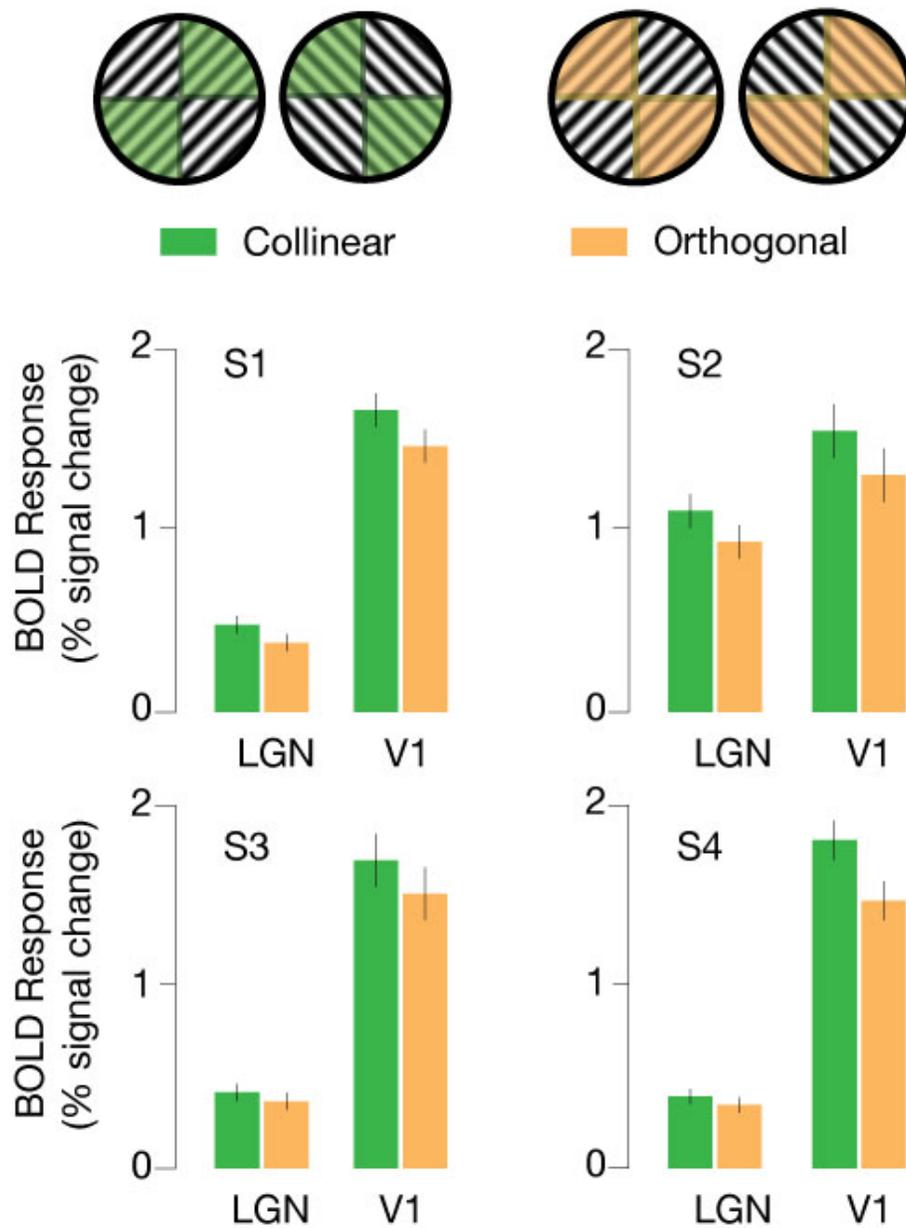
We further tested whether the same pattern of results were obtained for the two methods of ROI localization, by incorporating ROI localization method as a within-subjects factor into our ANOVA. We found no evidence of a reliable interaction effect between ROI localization method and any other factor, neither for decoding (all  $p$ 's > .1), nor for mean BOLD (all  $p$ 's > .6). In mean BOLD, there was a marginally significant trend of the factor of ROI localization method alone [ $F(1,47)=5.73$ ,  $p=0.06$ ].



Supplementary Figure 2

### Voxel-wise analysis of orientation biases

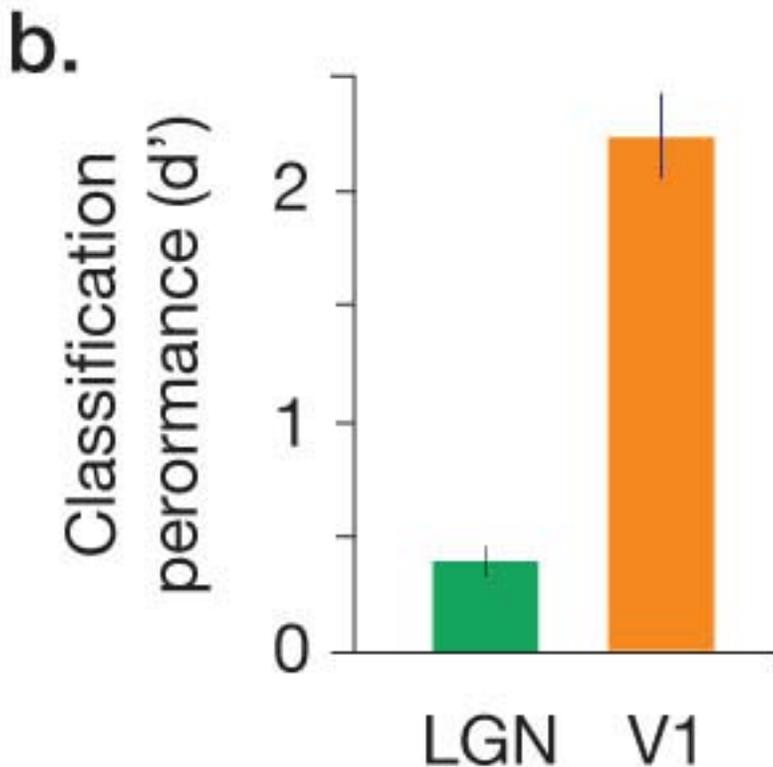
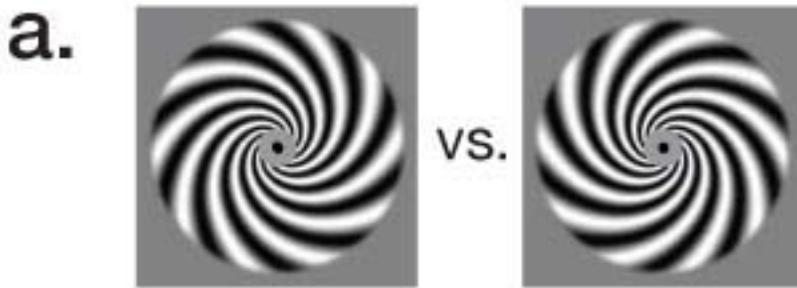
a) Voxel-wise analysis of orientation biases for attended and unattended gratings in the main experiment. To do so, we quantified the response difference between pairs of orientations using a Welch's t-test for every voxel in the LGN and V1 ROIs, within each condition. This metric allows us to directly quantify the strength of orientation biases in a given voxel, normalized by that voxel's reliability. Each color in the scatter plots (top) corresponds to orientation biases for an individual subject. Dashed lines correspond to the 95% cutoff for a null t-distribution. Evaluating the distribution of orientation preferences (t-values), we found that, across conditions, orientation preferences were distributed fairly normally (see border histograms for pooled t-values across subjects, per condition), although there appeared to be a bias towards vertical orientation responses compared to horizontal, in V1. By comparing the orientation biases per voxel for the attended condition against those in the unattended condition, we can also assess the degree to which voxels retain their orientation preferences between the two attentional conditions. Interestingly, the correlations in voxel preference between attended and unattended conditions were positive and significant for both oblique ( $\rho = 0.67$ ,  $p < .001$ ) and cardinal orientations ( $\rho = 0.74$ ,  $p < .001$ ) in V1, yet was only significant for cardinal orientations ( $\rho = 0.21$ ,  $p < .01$ ) in the LGN, and not for oblique orientations ( $\rho = -0.09$ ,  $p = 0.27$ ). This suggests that, aside from the oblique condition in the LGN, there was as significant correspondence in orientation preference for voxels between the attended and unattended conditions. In the LGN, however, oblique orientation representations appeared to emerge primarily in the presence of attention, which is consistent with our predictions based on the animal physiology (Vidyasagar & Urbas, 1982). b) To further quantify the strength of orientation information, we estimated the variance of the distributions of t-values within each condition. This variance provides a measure of the amount of orientation information present in the voxels on average: small variances indicate weak orientation biases, and large variances indicate strong orientation biases. Interestingly, this metric revealed a pattern of results that are qualitatively similar to our decoding results. First, the variance was substantially greater for V1 than LGN. Moreover, we observed an increase in this variance metric for both the oblique and cardinal conditions with attention in V1, whereas in the LGN the variance only changed with attention in the oblique condition. This additional set of analyses, which do not rely on multivariate pattern classification, provides converging evidence for our main result.



Supplementary Figure 3

**Data for individual subjects in the radial bias experiment**

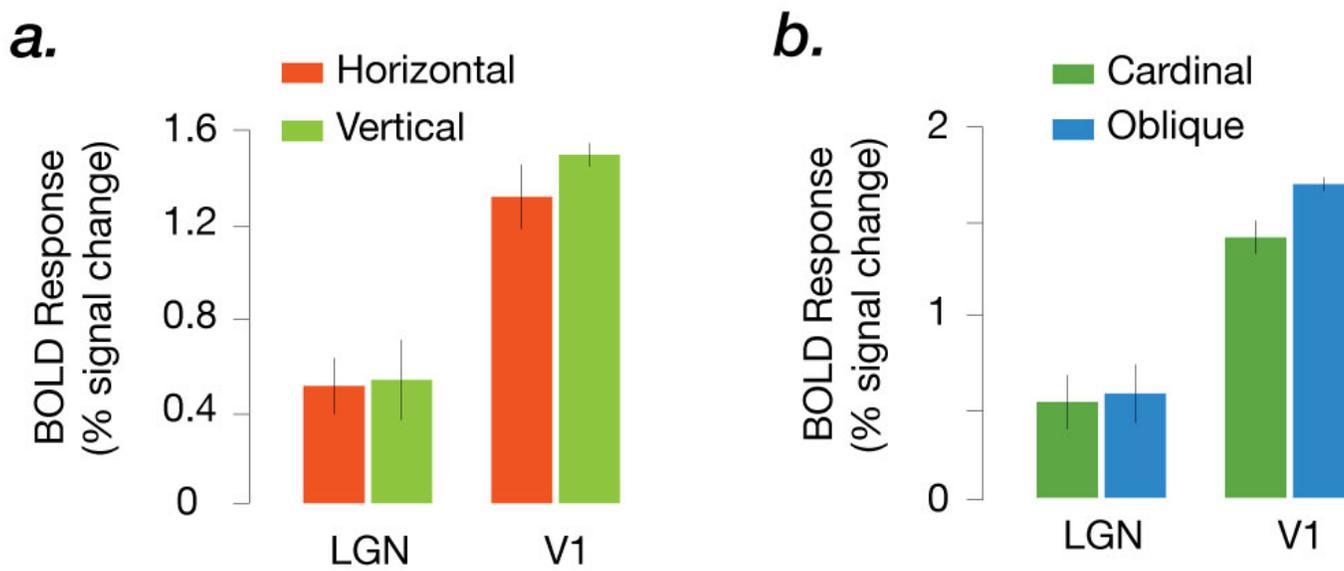
Mean BOLD responses were greater, both in LGN and V1, when the orientation of a stimulus matched the preferred retinotopic radial axis (collinear), compared to when it did not (orthogonal). These results suggest that orientation responses in the human LGN include a radial bias component, whereby orientation biases tend to be spatially clustered, in accordance to the polar angle of their retinotopic preference.



Supplementary Figure 4

**Decoding of spiral sense**

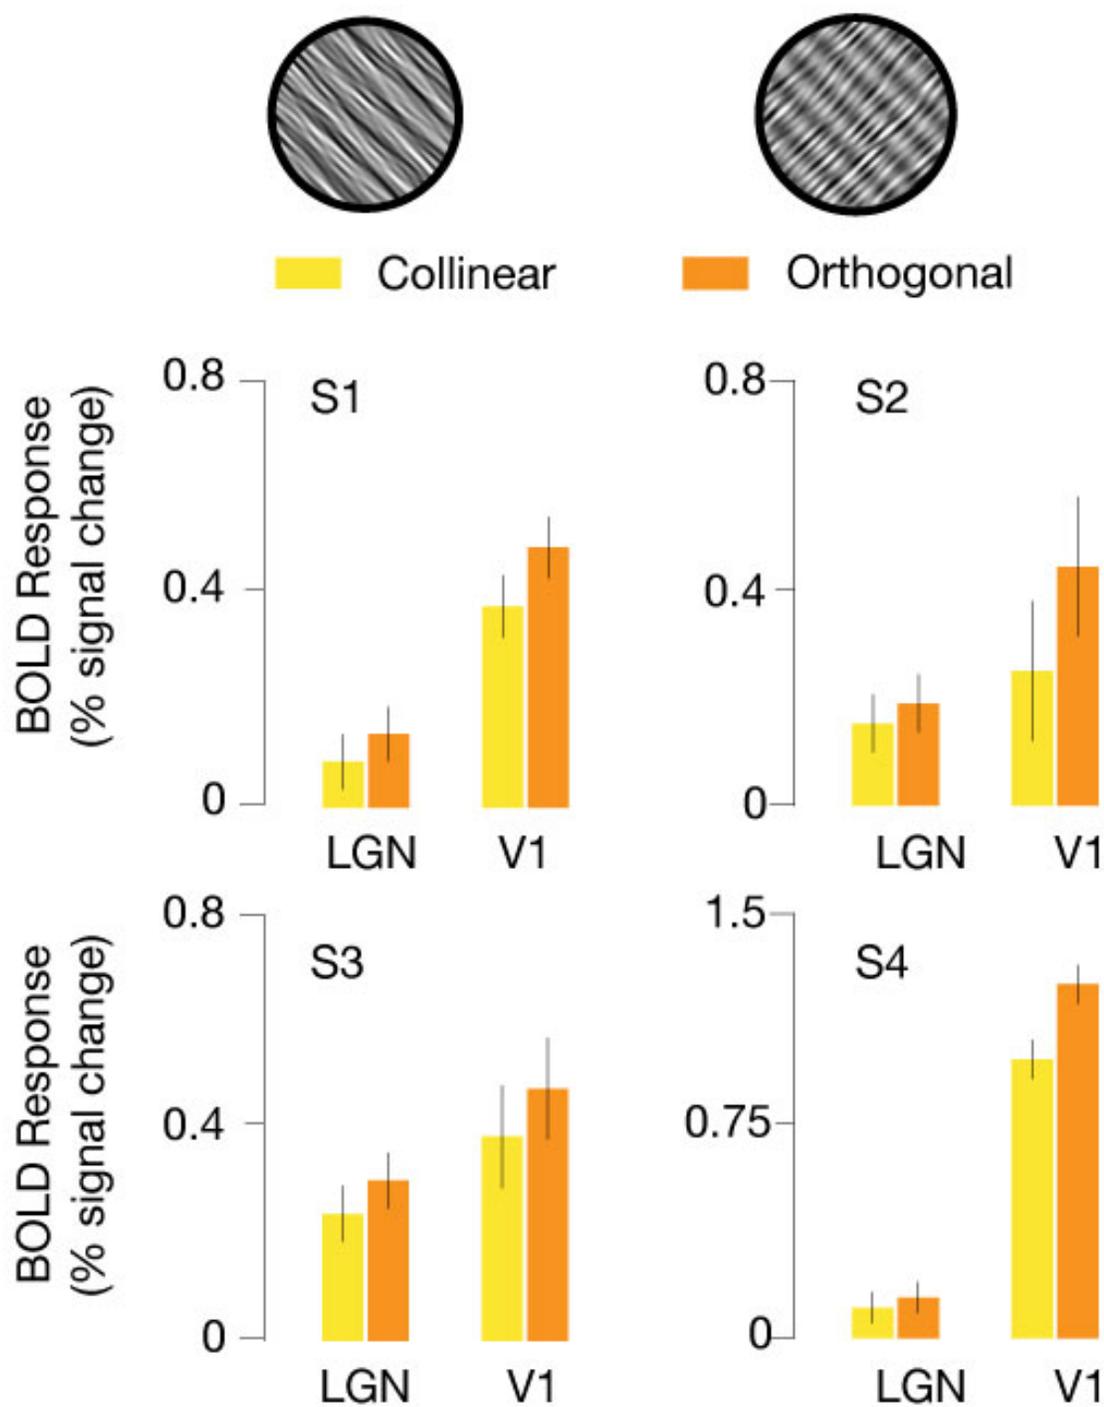
a) Example of stimuli used in a supplementary experiment, where we examined whether classification was possible for the orientation of logarithmic spiral gratings. Previous fMRI studies of the human visual cortex have shown that orientation-selective signals can be found at multiple spatial scales, ranging from the scale of cortical columns, to a coarse scale of >1cm, such as a retinotopically organized radial bias in which individual voxels exhibit a general preference for orientations that radiate away from the fovea (Alink, Krugliak, Walther, & Kriegeskorte, 2013; Freeman, Brouwer, Heeger, & Merriam, 2011; Sasaki et al., 2006). Spiral stimuli, however, can mitigate this radial bias (Mannion, McDonald, & Clifford, 2009), although other coarse-scale biases may be present (Freeman, Heeger, & Merriam, 2013). b) Mean classification performance of spiral orientation in LGN and V1 for 3 participants. As expected, classification performance was significant above chance for area V1,  $t(2)=12.17$ ,  $p<.05$ . Notably, orientation classification performance in the LGN, while lower than V1, was significantly above chance for each individual participant (all  $p$ 's<0.005 based on binomial test), and also for the group based on a one-sample t-test,  $t(2)=5.94$ ,  $p<.05$ .



Supplementary Figure 5

**Assessing response anisotropies between orientations.**

**a)** Comparison of mean BOLD responses to cardinal orientations in the LGN and V1. Results did not reveal any significant difference between vertical and horizontal orientations in the LGN ( $t(1,3) = 0.425$ ,  $p=0.349$ ), but a trend towards higher responses to vertical, compared to horizontal, in V1 ( $t(1,3) = 1.67$ ,  $p=0.097$ ). While our results did not reveal the preference for horizontal in the LGN that has been documented in some animal studies, this may be due to a variety of reasons, including that our ROI was defined using a large full-field localizer stimulus; while the orientation biases may exist between horizontal and vertical orientations, they are likely most prominent along their respective radial axes. A preference in BOLD response to vertical in the human visual cortex has been found in prior studies (Freeman et al, 2011; 2013). **b)** Comparison of mean BOLD response to cardinal and oblique orientations, measured within the same scan session. In LGN, we observed no significant difference in BOLD response between cardinal and obliques ( $t(1,3) = 1.34$ ,  $p>.1$ ), but did observe a significantly greater BOLD response to oblique than to cardinal orientations in V1 ( $t(1,3)=3.84$ ,  $p<.05$ ). Although some studies have reported greater mean BOLD activity for cardinals over obliques in early visual areas (Furmanski & Engel, 2000), the opposite pattern seen here has also been observed (Mannion, McDonald, & Clifford, 2010; Swisher et al., 2010), and interpreting the relationship between mean BOLD and orientation continues to be an active area of investigation.



Supplementary Figure 6

**Data for individual subjects in the Orientation-selective masking experiment**

BOLD responses were weaker, both in V1 and LGN, when the orientation bandpass-filtered noise and the sinusoidal grating shared the same orientation, as compared to when they were orthogonal.