



Published in final edited form as:

Nat Neurosci. 2015 April ; 18(4): 496–498. doi:10.1038/nn.3967.

Attention alters orientation processing in the human lateral geniculate nucleus

Sam Ling^{1,2}, Michael Pratte³, and Frank Tong³

¹Department of Psychological & Brain Sciences, and the Center for Computational Neuroscience & Neural Technology, Boston University ²Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, The Netherlands ³Department of Psychology, and the Vanderbilt Vision Research Center, Vanderbilt University

Abstract

Orientation selectivity is a cornerstone property of vision, commonly believed to emerge in the primary visual cortex (V1). Here, we demonstrate that reliable orientation information can be detected even earlier, in the human lateral geniculate nucleus (LGN), and that attentional feedback selectively alters these orientation responses. This attentional modulation may allow the visual system to modify incoming feature-specific signals at the earliest possible processing site.

What role does the human lateral geniculate nucleus (LGN) play in visual perception? Most models of vision treat the LGN as a passive relay station to V1^{1,2}, characterized by simple circular receptive fields. However, a growing body of neurophysiological evidence suggests that some neurons in the LGN possess receptive fields with elongated aspect ratios^{3–6} — a property that could support subcortical orientation-selective processing. Moreover, there is physiological evidence to suggest a plausible mechanism for orientation-specific feedback from V1 to the LGN^{7,8}. In particular, the distribution of orientation preferences in cat LGN has been shown to depend on whether V1 is lesioned or spared, with lesions leading to a diminished preference for oblique orientations but a preserved preference for cardinal orientations⁵. This result implies that feedback projections from V1 serve to strengthen the representation of oblique orientations in the LGN. In the present study, we present converging lines of evidence indicating that orientation signals can be detected in the human LGN. Furthermore, we demonstrate that attention selectively enhances representations of oblique orientations, revealing a potential consequence of attentional feedback on these orientation-selective responses.

Many fMRI studies have measured orientation-selective signals throughout the visual cortex, using multivariate pattern analysis to classify viewed orientation from voxel activity

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

Author Contributions

S.L., M.P. & F.T. conceived and designed the experiments. S.L. and M.P. collected the data. S.L. conducted the data analyses. S.L., M.P. & F.T. wrote the manuscript.

A supplementary methods checklist is available.

patterns^{9,10}. However, the presence of orientation-selective responses has not, to our knowledge, been demonstrated in the human LGN. fMRI BOLD activity was measured while participants were shown sinusoidal gratings, which were oriented either cardinally (0° or 90°) or obliquely (45° or 135°). The spatial phase of the gratings was randomized across presentations (1 Hz) within each 16-s stimulus block to ensure that pattern classification relied on orientation-selective information, rather than on retinotopic luminance differences⁹. Attention was manipulated by presenting a cue at the start of each block to indicate whether the participant should covertly attend to the oriented grating (*attended condition*) or to a sequence of letters that appeared in rapid succession at fixation (*unattended condition*). The grating task involved detecting and discriminating intermittent, near threshold changes in orientation, with magnitudes titrated to match for difficulty across conditions. The letter discrimination task involved reporting whenever a 'J' or 'K' appeared in the letter sequence. This central letter task is cognitively demanding, and effectively withdraws attention from the visual periphery.

Our first goal was to determine whether viewed orientation could be successfully classified from voxel activity patterns in the LGN. We localized the LGN region of interest in individual participants using a combination of functional^{11,12} and anatomical localization methods (Fig. 1a & Supplementary Fig. 1a). Our pattern analysis results indicated that reliable orientation information was present not only in V1, but in the human LGN, as well (Fig. 1b; all p 's < 0.01). Next, we asked whether the top-down effects of attentional feedback might influence orientation-selective responses in the LGN. To do so, we measured the extent to which attention improved orientation classification performance, separately for oblique and cardinal orientations, based on the hypothesis that cortical feedback might selectively bolster the representation of oblique orientations in the LGN⁵. We observed a significant interaction between the effects of attentional modulation, orientation, and visual area on classification performance, $F(4,40) = 5.51$, $p = 0.001$. In V1, attention increased the strength of orientation responses for both oblique and cardinal orientations (Fig. 1b, all p 's < .01) to a comparable extent, $t(5) = 0.39$, $p = 0.71$. In the LGN, however, the attentional effect was significantly greater for oblique orientations than for cardinal orientations (Fig. 1b; $t(5) = 4.76$, $p = 0.002$). Specifically, the withdrawal of attentional feedback led to substantially weaker orientation-selective activity patterns for oblique orientations in the LGN, but had no reliable effect on cardinal orientations (Fig. 1d). The qualitatively different pattern of results observed in the LGN and V1 suggests that a distinct type of feature-selective modulation is occurring in the LGN, leading to the selective enhancement of oblique orientations (Supplementary Figs. 1b & 2).

Do these attentional effects reflect changes in the pattern-specific component of these orientation responses, or might they reflect gross changes in mean BOLD activity? We compared the overall magnitude of the BOLD response for cardinal and oblique orientations. Attention led to enhanced responses in both the LGN and V1 (Fig. 1c, $F(4,40) = 20.19$, $p < .001$), but the magnitude of modulation was comparable for cardinal and oblique orientations in the LGN and V1, with no evidence of an interaction effect, $F(4,40) = 0.38$, $p = .56$. Taken together, our results indicate that attention enhances the overall BOLD response of the LGN regardless of stimulus orientation, but that the impact of attentional

feedback on feature selectivity is more nuanced: the degree to which attention affects LGN processing of a visual stimulus depends on stimulus orientation.

The amount of orientation information we observed in LGN activity patterns was modest, when compared to the high classification performance of area V1. This was to be expected due to several factors, including broader orientation selectivity of LGN neurons, poorer quality of fMRI signals from subcortical than cortical regions¹⁰, and the relatively smaller size of the LGN structure. Nevertheless, we observe reliable orientation-selective responses in the human LGN, and these may arise from several sources. For example, animal studies have suggested that ganglion cell receptive fields are not uniformly circular¹³, but instead exhibit modest orientation preferences organized at both fine¹⁴ and coarse spatial scales¹⁴. To determine whether one such coarse-scale orientation bias, known as the radial bias, exists in the human LGN, we conducted an additional experiment to examine whether LGN responses depend on the correspondence between stimulus orientation and retinotopic preference¹⁵. LGN voxels were localized based on their retinotopic preference for either of the two diagonal radial axes (Fig. 2a), and mean BOLD responses were significantly greater for full-field gratings that were collinear rather than orthogonal to a voxel's preferred radial axis (Fig. 2b, Supplementary Fig. 3, LGN: $t(3) = 6.649$, $p = .003$; V1: $t(3) = 7.536$; $p = .002$). These results indicate that the human LGN exhibits a coarse-scale preference for radial orientations, similar to what has been previously found in the human V1. In another experiment, we found that the orientation of logarithmic spiral gratings could also be decoded from LGN activity patterns, indicating the presence of other sources of orientation preference in human LGN, distinct from radial bias¹⁶ (Supplementary Fig. 4).

In a final experiment, we used univariate fMRI analyses to test whether the LGN is sensitive to orientation-tuned masking. By combining target and mask images that have collinear or orthogonal orientations, one can test for nonlinear orientation-selective responses to the linearly summed image pairs^{17,18}. We found that BOLD responses in the LGN and V1 exhibit a predicted dependence on orientation: responses were weaker when target and mask were collinear rather than orthogonal (Fig. 2c & d, LGN: $t(3) = 7.44$, $p = .002$; V1: $t(3) = 4.11$; $p = .013$). These differential responses could not result from LGN detectors that lacked orientation tuning, providing further evidence that orientation responses are present in the human LGN (Supplementary Fig. 6).

Recent studies demonstrate that top-down attention can impact the neural responsivity of the LGN^{11,19}, but we are in the nascent stages of understanding of how it might affect feature-selective processing²⁰. Our study provides converging evidence, using both multivariate and univariate fMRI approaches, for orientation-selective responses in the human LGN. These orientation-selective responses exhibit some degree of coarse organization coincident with retinotopic preference, and they appear to exhibit orientation-dependent interactions. Taken together, these results suggest that the LGN may play an underappreciated role in processing orientation information, a process that has traditionally been attributed to the visual cortex. Moreover, we found that attention altered these orientation-selective responses suggesting a remarkably early locus by which top-down processes act upon feature-selective responses. Even modest levels of attentional modulation at this stage could potentially lead to larger

effects downstream, suggesting that this early modulation may have an important role in visual perception.

Methods

Observers

Nine healthy adult volunteers (aged 22–41 years, 3 female), with normal or corrected-to-normal vision, participated in the study. All experiments were matched for sex (aside from the spirals study, which had 2 females and 1 male). All subjects gave informed written consent. The study was approved by the Vanderbilt University Institutional Review Board. In each study, participants completed 13–22 fMRI runs (~5 min duration each) per scanning session. A power analysis indicated that 6 subjects would be sufficient to detect the predicted decoding and attention effects. Indeed, this sample size is consistent with previous fMRI decoding studies^{21–23}.

Apparatus

The stimuli were generated using Matlab and the Psychophysics Toolbox^{24,25}, and were displayed on a rear-projection screen using a gamma-corrected Eiki LC-X60 LCD projector with a Navitar zoom lens, which participants viewed through a front surface mirror. To minimize head motion, participants used a customized bite bar and padding.

MRI acquisition

Most functional and anatomical data were acquired on a Philips 3T Intera Achieva MRI scanner, equipped with an 8-channel head coil. Each MRI session lasted 2.5 hours, during which we acquired: i) a T1-weighted 3D anatomical scan of the entire brain (1 mm isotropic), ii) a T2-weighted structural in-plane for EPI alignment, iii) 2–4 functional runs to identify retinotopic regions in the LGN and visual cortex that correspond to the stimulus location, and iv) 13–22 fMRI runs to measure BOLD activity during the experimental task. BOLD activity was measured using gradient-echo T2*-weighted echoplanar imaging; twenty slices were acquired axially, with through-plane coverage of the thalamus and the occipital pole, and a voxel size of $2 \times 2 \times 2$ mm (TR 2 s, TE 35 ms; flip angle 79° ; FOV 192×192 mm). Per subject, the study consisted of 2 scan sessions for the attention study, 1 session for the spirals study, 1 session for the radial bias study, and 1 session for the tuned masking study.

Proton density-weighted (PD) images, which were used to anatomically delineate the LGN, were acquired in a separate scanning session on a Philips 7T Achieva MRI scanner, equipped with a 32-channel head coil. Each PD scan lasted 1 hour, during which T1-weighted structural in-plane images was acquired for alignment in addition to 18–20 PD volume acquisitions. PD images were acquired with whole-brain coverage, and a voxel size of $1 \times 1 \times 1$ mm (TR 4.3 s, TE 1.9 ms).

Stimuli and design

Attention experiment—Participants were instructed to maintain fixation on a small circle (diameter: 0.1°) at the center of the display throughout the course of the experiment.

Retinotopic regions of interest (ROIs) were identified in the visual cortex and the LGN based on 2 functional localizer scans, in which subjects viewed a flickering checkerboard of the same size and location as the stimuli used in the main experiment (presented in a circular aperture centered on fixation, with an outer diameter of 16° , and an inner diameter of 0.7°). Each localizer run started with 12 s of fixation, followed by 12 s of a flickering checkerboard stimulus (100% contrast; counter-phase reversing at 6 Hz). This off/on stimulus cycle was repeated 8 times in each run, and ended with a final 12 s of fixation. A given run took, in total, 204 seconds.

For the main experiment, observers viewed gratings presented at 0° , 45° , 90° or 135° . The cardinal (0° =vertical, 90° =horizontal) and oblique (45° , 135°) orientations were presented in different scan sessions, and the ordering of scans was counterbalanced across subjects (experimenter was not blind to condition), such that 3 subjects were scanned with the oblique orientations first, and the other three were scanned with the cardinal orientations first. To ensure that the classifier relied on orientation-selective information, rather than retinotopic luminance differences, the spatial phase of pairs of gratings was randomized (500 ms on, 500 ms off) within each 16-s stimulus block⁹. To further rule out the possibility that fine-scale retinotopic differences in luminance might contribute to orientation decoding, we attempted to decode stimulus orientation based on the pixel-wise average of counter-phased stimuli, across each block of trials. Decoding performance of these images was at chance, confirming that above-chance decoding of LGN activity patterns must arise from orientation-selective responses⁹. Each run started with 16 s of fixation, followed by 16 s of an oriented grating. This was repeated 16 times in each run, and ended with a final 16 s of fixation. Each run took 288 seconds to complete. Each fMRI run included an equal number of stimulus blocks for each orientation, presented in a randomly generated order.

To manipulate attention towards or away from the oriented gratings, a series of letters was presented at central fixation throughout the experiment. A cue appeared at the beginning of each block to indicate whether the participant should attend to the oriented grating (attended condition) or to the sequence of letters at fixation (unattended condition). The grating task involved discriminating near-threshold changes in orientation, and difficulty was titrated to be comparable (~90% accuracy) across conditions and participants. For the letter discrimination task, observers were instructed to report whenever a 'J' or 'K' appeared in the letter sequence (5 letters/s). The attentional task alternated between blocks within each run.

Spiral grating experiment—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 $>1\text{cm}$, such as a retinotopically organized radial bias in which individual voxels exhibit a general preference for orientations that radiate away from the fovea^{15,26,27}. Spiral stimuli, however, can mitigate this radial bias¹⁶, although other coarse-scale biases may be present²⁸. Three participants were shown annular logarithmic spiral gratings in a block design (Supplementary Figure 2; $1\text{--}9^\circ$ eccentricity, phase-randomized: 500 ms on, 500 ms off, 10 cycles/rotation), at either $\pm 45^\circ$ pitch. Each run started with 16 s of fixation, followed by 16 s of an oriented grating. This was repeated 16 times in each run, and ended with a final 16 s of fixation. Each run took 288 seconds to

complete. Each fMRI run included an equal number of stimulus blocks for each spiral pitch, presented in a randomly generated order. Classification performance for the pitch of the spirals was evaluated to test for orientation information in LGN and V1 that does not correspond strictly to the radial bias.

Radial bias experiment—To assess the relationship between orientation preference and topography within the LGN and V1, we conducted a separate experiment to explicitly test for radial orientation biases in the LGN. Using a paradigm adopted from Sasaki et al (2006), we examined whether BOLD responses for voxels in the LGN and V1 depended on whether the orientation of a stimulus fell along the axis corresponding to the voxel's retinotopic polar angle preference, or was orthogonal to the radial axis (Figure 2a). On functional localizer runs, flickering checkerboards were presented in diagonally opposing quadrants (Figure 2a) to identify voxels with retinotopic preference that corresponded better with either the 45° or 135° radial axis. Each localizer run started with 12 s of fixation, followed by 12 s of a wedge-pair stimulus (100% contrast; counter-phase reversing at 6 Hz). This off/on stimulus cycle was repeated 10 times in each run, alternating between the 45° and 135° wedge-pairs across successive cycles, and ended with a final 12 s of fixation. Each localizer run lasted 252 seconds.

In the remaining runs, full-field gratings oriented at 0°, 45°, 90° or 135° (100% contrast) were shown to assess whether BOLD responses depended on an interaction between retinotopic preference and orientation. The spatial phase of the grating was phase-reversed (500 ms on, 500 ms off) within each 12-s stimulus block⁹. Each run started with 12 s of fixation. This was followed by repeated cycles consisting of 12 s of an oriented grating, and 12 s of subsequent blank fixation. This off/on stimulus cycle was repeated 10 times in each run, and ended with a final 12 s of fixation. Each run lasted 312 seconds. Each fMRI run included an equal number of stimulus blocks for each orientation, presented in a randomly generated order. To maintain fixation, a series of letters was presented at central fixation throughout the experiment, and observers were instructed to report whenever a 'J' or 'K' appeared in the letter sequence (5 letters/s).

The inclusion of 0° and 90° orientations, while not relevant in this design for the investigation of the radial bias, allowed us to also evaluate the difference in BOLD response between oblique (45° and 135°) and cardinal orientations (0° and 90°) within the same scanning session (Supplementary Figure 4).

Orientation tuned masking experiment—To further assess the evidence for orientation responses in the human LGN, we combined a critical-band noise-masking technique^{18,29–31} and univariate fMRI analyses to test for sensitivity to orientation-specific masking. LGN and V1 ROI's were acquired using a localizer paradigm identical to the attention experiment. Noise masking involves embedding an oriented signal within orientation bandpass-filtered noise. Under these conditions, a population of orientation-selective units will respond more weakly as the orientation content of the signal and noise components becomes more similar, as the stimuli will activate a common set of orientation units rather than two distinct sub-populations. Simulations of neural responses to such stimuli, using both a V1 simple cell Gabor model with half-wave rectification^{32,33} and

population coding models³⁴ confirmed that a population of orientation detectors will respond more weakly as the signal and noise orientations become more similar. In contrast, a population of non-oriented center-surround units, based on a difference of Gaussians model³⁵, yielded a flat pattern of responses as a function of the orientation difference between signal and noise components. Here, we apply this technique to mean BOLD responses in both V1 and LGN, to assess whether the LGN reveals effects of orientation-dependent masking.

Participants viewed a sinusoidal grating oriented at either 45° or 135° (40% contrast, 2 cycles/°, counter-phasing at 4 Hz), embedded in orientation bandpass-filtered noise (40% contrast, orientation bandwidth of 5°). The center frequency of this noise band was either identical to the grating orientation (collinear) or rotated 90° from the grating orientation (orthogonal). Stimuli were presented using a block design to measure mean BOLD responses to each of the four possible combinations of grating orientation and noise orientation (45°/45°, 45°/135°, 135°/45°, and 135°/135°). Each run started with 12 s of fixation. This was followed by repeated cycles consisting of 12 s of an oriented grating, and 12 s of subsequent blank fixation. This off/on stimulus cycle was repeated 16 times in each run, and ended with a final 12 s of fixation. Each run lasted 408 s. The order of conditions was pseudo-randomized in a given scan. To encourage fixation, a series of letters was presented at central fixation throughout the experiment, and observers were instructed to report whenever a 'J' or 'K' appeared in the letter sequence (5 letters/s). Each fMRI run included an equal number of stimulus blocks for each orientation, presented in a random order.

fMRI analyses

The functional data were preprocessed using standard motion-correction procedures³⁶. For each run, the raw MRI time series for every voxel was converted to units of percent signal change by dividing each voxel by its mean intensity across that run. To functionally localize ROIs, the localizer time series were fit with a general linear model, which assumed a temporally shifted double-gamma function as a model of the hemodynamic response function. This localizer was combined with standard retinotopic mapping procedures on data obtained in a separate scan session to identify the 150 most responsive voxels within area V1, in the native space for each participant. Contiguous voxels in the thalamus that responded to the localizer stimulus at a threshold of $p < 0.01$ defined the LGN functional region of interest (ROI)^{11,37–39}. To anatomically localize the LGN, the PD images were motion-corrected and averaged across scans, and the LGN was identified as the contiguous voxels within the thalamus that appeared as a higher intensity, teardrop-shaped structure medial and superior to the hippocampus, and lateral to the pulvinar⁴⁰, often with a gap corresponding to the different density of the white matter sitting between the LGN and the pulvinar. Analyses were conducted using 25 ± 7 (mean \pm s.d.) voxels across both left and right LGN. The selection of voxels from this anatomical localizer was conservatively lateralized, to avoid including the adjacent lateral inferior pulvinar. Pattern classification in the LGN was performed on voxels that corresponded to the intersection of the functionally defined LGN ROI and the structurally defined ROI. On average, 61% of proton density-defined LGN overlapped with the functionally defined LGN. Note also that the pattern of

results when using only the functional localizer look quite similar to the restricted ROI. Thus, it is unlikely that our conservative voxel selection biased the analysis. (Supplementary Figure 1). In the radial bias experiment, a GLM was fit to the localizer data to generate two ROI's, distinguishing between voxels that responded to the 45° or 135° wedge-pairs.

For the attention experiment, patterns of activity for each stimulus block were obtained by temporally averaging the BOLD activity across each block for every voxel in an ROI, after time shifting by 4 seconds to account for hemodynamic lag. Multivoxel pattern analysis was performed using linear support vector machines (SVM)⁹. Classifiers were trained using a 'leave-one-run-out' cross-validation scheme, with performance assessed by testing each left-out run and averaging the resulting classification accuracies over runs. The penalty parameter, C, was tuned within every training set using a cross-validation procedure, which ensured an unbiased selection of C⁴¹. Classification results presented for the attention experiment represent performance *within* the cardinal and oblique orientation pairs. For all analyses, classification accuracy was converted to d' ⁴². The interaction was tested with repeated-measures ANOVA. A power analysis revealed that 6 subjects would be sufficient to detect the predicted decoding and attention effects. Indeed, this sample size is consistent with previous fMRI decoding studies^{21–23}.

In addition to the reported t-tests showing that decoding was reliably above chance at the group level, we performed a subject-wise permutation analysis in which chance-level decoding performance was calculated by shuffling the orientation labels 1000 times, per subject, experimental condition, and ROI. In V1, decoding performance fell outside of the bounds of the null distribution's 95% confidence interval in all instances (24/24). In the LGN, decoding performance fell outside of the bounds of the null in 12/12 instances (i.e., every subject and orientation condition) in the attended condition, and 9/12 instances for unattended.

In the attention experiment, the block design used for classification did not have any blank fixation periods between blocks (to maximize the number of samples for classification). Because our measure of % BOLD signal change is normalized to the entire time series, the resulting response estimates for attended and unattended conditions are positive and negative respectively, rather than all being above zero. In the Radial Bias and Orientation-Tuned Masking experiments, BOLD responses were estimated per condition by fitting the time series with a GLM. However, in these experiments there were epochs of blank fixation (12 s) between each stimulus presentation, thereby lead to positive values for % signal change BOLD responses across all conditions.

Eye-position monitoring

Participants were instructed to maintain fixation on a central fixation point throughout each fMRI run. For the experiment involving spiral stimuli, the participant's eye position was monitored using an MRI-compatible ASL EYE-TRAC eye-tracking system. We applied pattern classifiers to these data to evaluate whether any reliable information about the viewed orientation could be decoded from eye-position signals²². Analyses of eye movement data during a subset of scans revealed that our ability to decode orientation from eye position was at chance level, $t(3)=0.34$, $p=0.75$. We also conducted additional high-

resolution measurements of eye position (500 Hz, SR-Research Eyelink II) outside of the MRI scanner for 5 of the subjects that ran in the study, using a visual stimulation paradigm identical to that used in the main attention experiment. We found that when we tried to decode the orientation of a viewed grating based solely on the high-resolution eye tracking data, performance was at chance level, $t(4)=1.50$, $p>0.2$. These findings replicate other published results from our lab, in which we attempted to decode stimulus orientation from eye position^{21,22,43}. Taken together, these results indicate that eye movements are unlikely to have contributed to our ability to classify orientation based on fMRI activity patterns.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank V. Casagrande, J. Schall, & R. Blake for valuable comments and discussion. Supported by US National Institutes of Health grant R01 EY017082 to FT, R01 EB000461 (Director Dr. Gore), and NIH Fellowship F32-EY022569 to MP.

References

1. Hubel DH, Wiesel TN. Receptive fields of single neurones in the cat's striate cortex. *J Physiol.* 1959; 148:574–591. [PubMed: 14403679]
2. Ferster D, Miller KD. Neural mechanisms of orientation selectivity in the visual cortex. *Annu Rev Neurosci.* 2000; 23:441–471. [PubMed: 10845071]
3. Piscopo DM, El-Danaf RN, Huberman AD, Niell CM. Diverse visual features encoded in mouse lateral geniculate nucleus. *Journal of Neuroscience.* 2013; 33:4642–4656. [PubMed: 23486939]
4. Xu X, Ichida J, Shostak Y, Bonds AB, Casagrande VA. Are primate lateral geniculate nucleus (LGN) cells really sensitive to orientation or direction? *Vis Neurosci.* 2002; 19:97–108. [PubMed: 12180863]
5. Vidyasagar TR, Urbas JV. Orientation sensitivity of cat LGN neurones with and without inputs from visual cortical areas 17 and 18. *Exp Brain Res.* 1982; 46:157–169. [PubMed: 7095028]
6. Cheong SK, Tailby C, Solomon SG, Martin PR. Cortical-like receptive fields in the lateral geniculate nucleus of marmoset monkeys. *Journal of Neuroscience.* 2013; 33:6864–6876. [PubMed: 23595745]
7. Wang W, Jones HE, Andolina IM, Salt TE, Sillito AM. Functional alignment of feedback effects from visual cortex to thalamus. *Nat Neurosci.* 2006; 9:1330–1336. [PubMed: 16980966]
8. Murphy PC, Duckett SG, Sillito AM. Feedback connections to the lateral geniculate nucleus and cortical response properties. *Science.* 1999; 286:1552–1554. [PubMed: 10567260]
9. Kamitani Y, Tong F. Decoding the visual and subjective contents of the human brain. *Nat Neurosci.* 2005; 8:679–685. [PubMed: 15852014]
10. Fischer J, Whitney D. Attention gates visual coding in the human pulvinar. *Nat Commun.* 2012; 3(1051)
11. O'Connor DH, Fukui MM, Pinsk MA, Kastner S. Attention modulates responses in the human lateral geniculate nucleus. *Nat Neurosci.* 2002; 5:1203–1209. [PubMed: 12379861]
12. Schneider KA, Kastner S. Effects of sustained spatial attention in the human lateral geniculate nucleus and superior colliculus. *Journal of Neuroscience.* 2009; 29:1784–1795. [PubMed: 19211885]
13. Schall JD, Perry VH, Leventhal AG. Retinal ganglion cell dendritic fields in old-world monkeys are oriented radially. *Brain Res.* 1986; 368:18–23. [PubMed: 3955359]

14. Shou TD, Leventhal AG. Organized arrangement of orientation-sensitive relay cells in the cat's dorsal lateral geniculate nucleus. *Journal of Neuroscience*. 1989; 9:4287–4302. [PubMed: 2593002]
15. Sasaki Y, et al. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron*. 2006; 51:661–670. [PubMed: 16950163]
16. Mannion DJ, McDonald JS, Clifford CW. Discrimination of the local orientation structure of spiral Glass patterns early in human visual cortex. *Neuroimage*. 2009; 46:511–515. [PubMed: 19385017]
17. Ni AM, Ray S, Maunsell JH. Tuned normalization explains the size of attention modulations. *Neuron*. 2012; 73:803–813. [PubMed: 22365552]
18. Ling S, Pearson J, Blake R. Dissociation of neural mechanisms underlying orientation processing in humans. *Current Biology*. 2009; 19:1458–1462. [PubMed: 19682905]
19. McAlonan K, Cavanaugh J, Wurtz RH. Guarding the gateway to cortex with attention in visual thalamus. *Nature*. 2008; 456:391–394. [PubMed: 18849967]
20. Schneider KA. Subcortical mechanisms of feature-based attention. *Journal of Neuroscience*. 2011; 31:8643–8653. [PubMed: 21653868]

Methods References

21. Harrison SA, Tong F. Decoding reveals the contents of visual working memory in early visual areas. *Nature*. 2009; 458:632–635. [PubMed: 19225460]
22. Jehee JF, Brady DK, Tong F. Attention improves encoding of task-relevant features in the human visual cortex. *Journal of Neuroscience*. 2011; 31:8210–8219. [PubMed: 21632942]
23. Pratte MS, Ling S, Swisher JD, Tong F. How attention extracts objects from noise. *J Neurophysiol*. 2013; 110:1346–1356. [PubMed: 23803331]
24. Brainard DH. The Psychophysics Toolbox. *Spat Vis*. 1997; 10:433–436. [PubMed: 9176952]
25. Pelli DG. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis*. 1997; 10:437–442. [PubMed: 9176953]
26. Freeman J, Brouwer GJ, Heeger DJ, Merriam EP. Orientation decoding depends on maps, not columns. *Journal of Neuroscience*. 2011; 31:4792–4804. [PubMed: 21451017]
27. Alink A, Krugliak A, Walther A, Kriegeskorte N. fMRI orientation decoding in V1 does not require global maps or globally coherent orientation stimuli. *Front Psychol*. 2013; 4:493. [PubMed: 23964251]
28. Freeman J, Heeger DJ, Merriam EP. Coarse-Scale Biases for Spirals and Orientation in Human Visual Cortex. *Journal of Neuroscience*. 2013; 33:19695–19703. [PubMed: 24336733]
29. Stromeyer CF 3rd, Julesz B. Spatial-frequency masking in vision: critical bands and spread of masking. *J Opt Soc Am*. 1972; 62:1221–1232. [PubMed: 4643008]
30. Tsai JJ, Wade AR, Norcia AM. Dynamics of normalization underlying masking in human visual cortex. *Journal of Neuroscience*. 2012; 32:2783–2789. [PubMed: 22357861]
31. Patterson RD. Auditory filter shapes derived with noise stimuli. *J Acoust Soc Am*. 1976; 59:640–654. [PubMed: 1254791]
32. Serre T, Wolf L, Bileschi S, Riesenhuber M, Poggio T. Robust object recognition with cortex-like mechanisms. *IEEE Trans Pattern Anal Mach Intell*. 2007; 29:411–426. [PubMed: 17224612]
33. Grigorescu SE, Petkov N, Kruizinga P. Comparison of texture features based on Gabor filters. *IEEE Trans Image Process*. 2002; 11:1160–1167. [PubMed: 18249688]
34. Averbeck BB, Latham PE, Pouget A. Neural correlations, population coding and computation. *Nat Rev Neurosci*. 2006; 7:358–366. [PubMed: 16760916]
35. Enroth-Cugell C, Robson JG. Functional characteristics and diversity of cat retinal ganglion cells. Basic characteristics and quantitative description. *Invest Ophthalmol Vis Sci*. 1984; 25:250–267. [PubMed: 6698746]
36. Nestares O, Heeger DJ. Robust multiresolution alignment of MRI brain volumes. *Magn Reson Med*. 2000; 43:705–715. [PubMed: 10800036]
37. Kastner S, et al. Functional imaging of the human lateral geniculate nucleus and pulvinar. *J Neurophysiol*. 2004; 91:438–448. [PubMed: 13679404]

38. Kastner S, Schneider KA, Wunderlich K. Beyond a relay nucleus: neuroimaging views on the human LGN. *Prog Brain Res.* 2006; 155:125–143. [PubMed: 17027384]
39. Wunderlich K, Schneider KA, Kastner S. Neural correlates of binocular rivalry in the human lateral geniculate nucleus. *Nat Neurosci.* 2005; 8:1595–1602. [PubMed: 16234812]
40. Devlin JT, et al. Reliable identification of the auditory thalamus using multi-modal structural analyses. *Neuroimage.* 2006; 30:1112–1120. [PubMed: 16473021]
41. Ku SP, Gretton A, Macke J, Logothetis NK. Comparison of pattern recognition methods in classifying high-resolution BOLD signals obtained at high magnetic field in monkeys. *Magn Reson Imaging.* 2008; 26:1007–1014. [PubMed: 18691999]
42. Green, DM.; Swets, JA. *Signal detection theory and psychophysics.* Wiley: 1966.
43. Jehee JF, Ling S, Swisher JD, van Bergen RS, Tong F. Perceptual learning selectively refines orientation representations in early visual cortex. *Journal of Neuroscience.* 2012; 32:16747–16753a. [PubMed: 23175828]

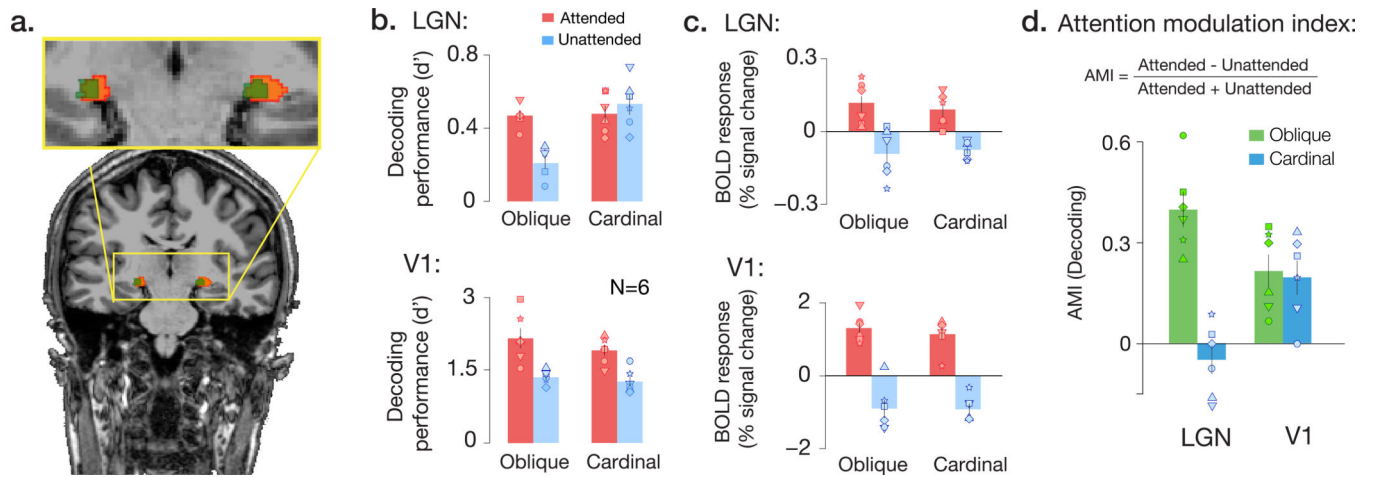


Figure 1.

Attention selectively augments orientation representation in the human LGN. **a**) LGN region of interest in a representative subject, identified by the intersection (green) of functional localizers (orange) and proton density-weighted structural imaging (blue). **b**) Accuracy of orientation decoding (d' units) for multivariate activity patterns in LGN and V1, for attended and unattended gratings. Attention led to enhanced orientation-selective responses in V1 for all orientation conditions, whereas in the LGN only oblique orientations were modulated by attention. Oblique and cardinal orientation pairs were tested on different scan sessions. Individual points correspond to individual subjects. A non-parametric permutation test confirmed that V1 decoding performance fell outside of the 95% CI bounds of the null distribution in all conditions for every individual participant (24/24 cases), and that LGN decoding performance fell outside of the bounds of the null in 12/12 cases in the attended condition, and 9/12 cases for the unattended condition. **c**) Attention had comparable effects on mean BOLD activity across orientations, in both the LGN and V1. BOLD response was normalized by the mean intensity across the time series for each run. **d**) Attentional modulation indices (AMI) for decoding performance of oblique (green) and cardinal (blue) orientations, in LGN and V1. Higher positive values indicate larger effects of attention. Whereas attentional modulation was comparable for cardinal and oblique orientations in area V1, attention selectively modulated responses to oblique orientations in the LGN. Error bars denote ± 1 s.e.m.

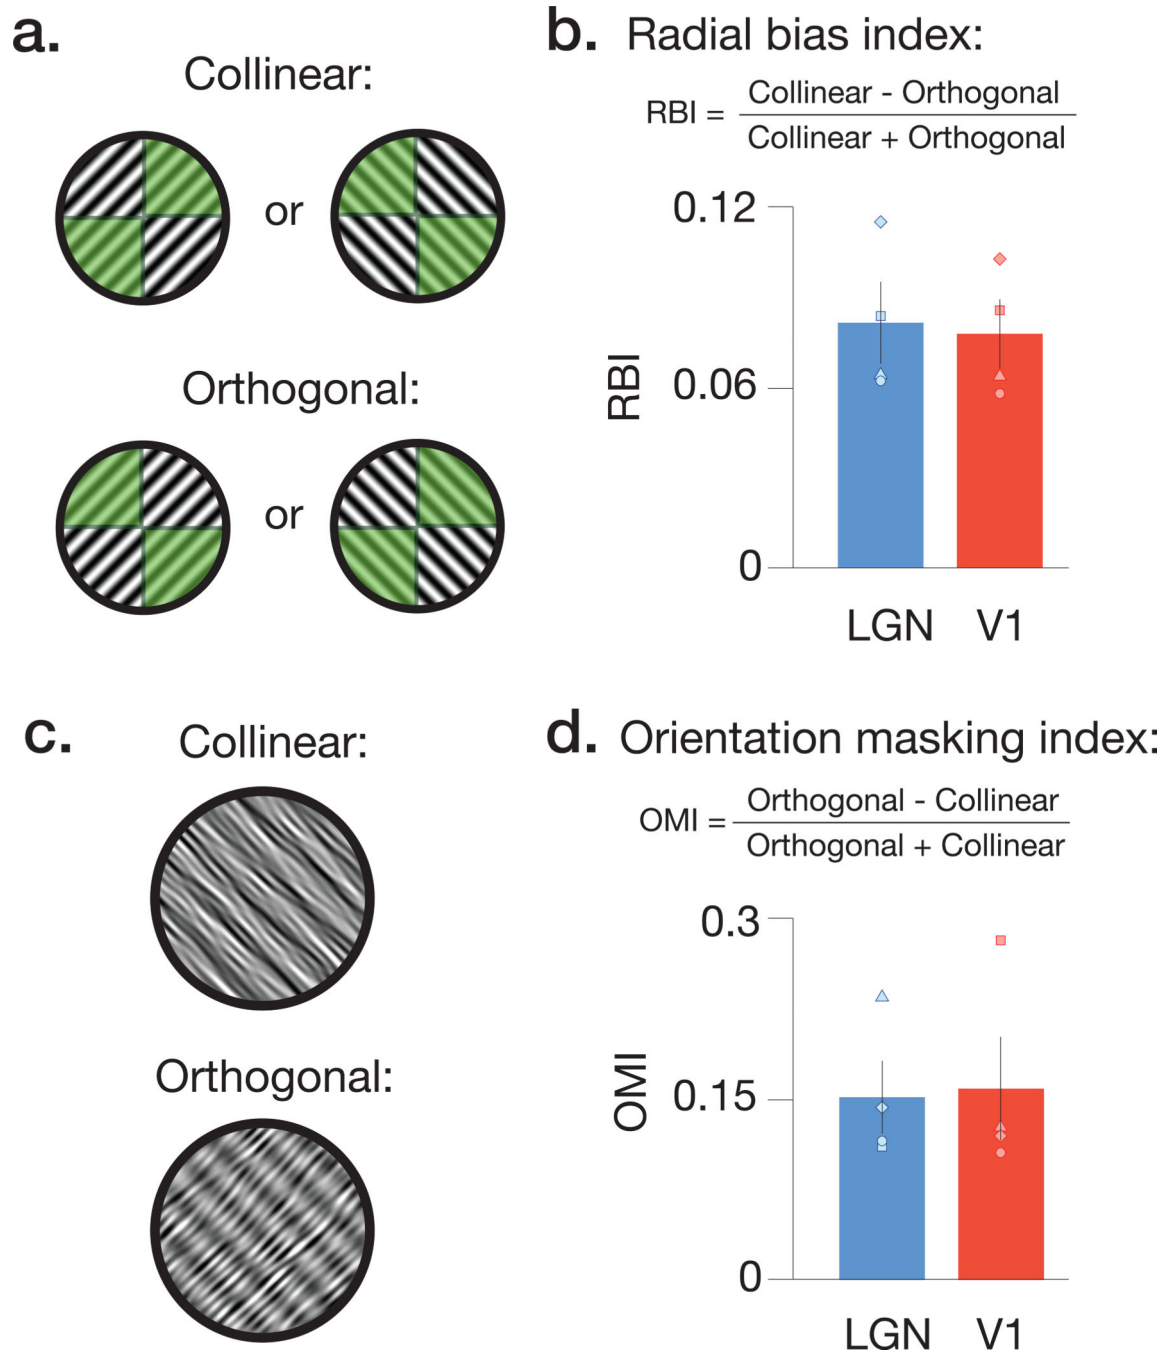


Figure 2.

Retinotopic preference for radial orientation and effects of orientation-specific masking in the human LGN and V1. **a)** Schematic of configurations used to test for radial bias. Functional localizers were used to determine voxels with retinotopic preference that fell along either the 45° or 135° axis (depicted by green shaded areas with diagonally opposing quadrants). We then presented a series of full-field gratings, which were oriented 45° or 135°. The relationship between localizer configuration and stimulus orientation determined the *Collinear* and *Orthogonal* conditions. **b)** Radial bias indices for mean BOLD responses

in both LGN and V1. Higher positive values indicate larger radial biases. The results reveal that, in both V1 and LGN, the strength of the response depended critically on the match between the orientation of a stimulus, and the retinotopic preference of a region of interest. **c)** Illustration of stimuli used to test for orientation-tuned masking in LGN and V1. Stimuli were composed of linear sinusoidal gratings summed with orientation bandpass filtered noise. The noise orientation and grating were configured either collinear or orthogonal to each other. **d)** Orientation masking indices for mean BOLD responses in both LGN and V1 revealed that collinear stimuli had stronger suppression (lower BOLD response) than orthogonal stimuli. Error bars denote ± 1 s.e.m.