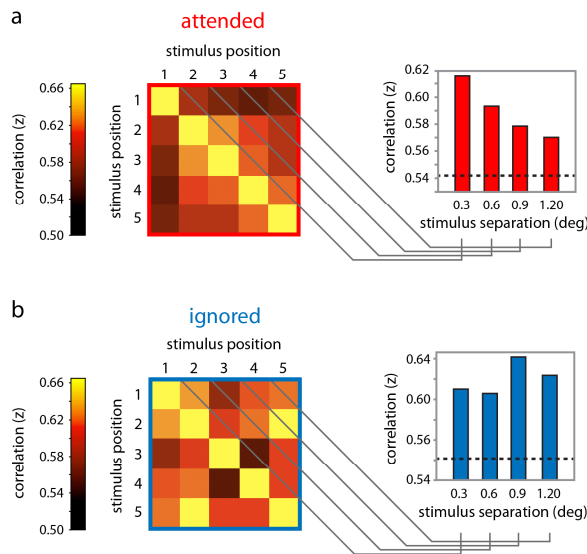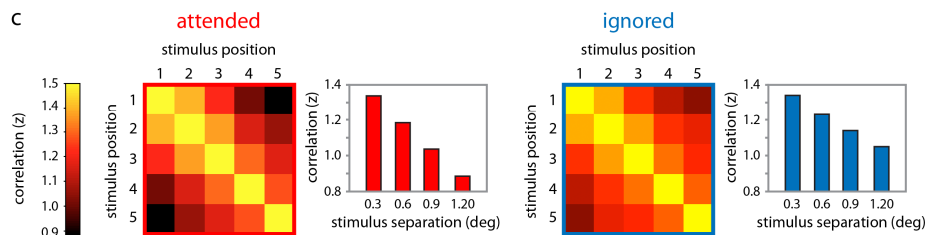Pulvinar:

Supplementary Figure S1. Correlation matrices for the pulvinar and V1.
a) Correlations among the maps for the five attended stimulus positions within the pulvinar, collapsed across hemispheres and subjects. Red bars show the mean correlation for each of the four stimulus separations, and the dashed line within the bar plot shows the mean of the bootstrapped null distribution (±1 sd of the distribution fall within the thickness of the dashed line; see Supplemental Methods). The null distribution reflects a baseline level of correlation among the five maps that cannot be attributed to stimulus-related activity, and might instead reflect factors such as partial voluming, vascular structure, etc. The correlation at each of the four stimulus separations was significantly greater than baseline (all $p < 0.001$, evaluated as the proportion of the bootstrapped null distribution that was greater than the actual correlation between attended maps; bootstrapping based on 80 correlation values per subject). There was a significant negative trend in correlation with increasing stimulus separation (Fisher $z = -0.46$; $p < 0.001$; based on 50 points in a position discrimination plot constructed as in Fig. 5, but now including data at stimulus separations of 0.9 and 1.2 deg). Correlations here were measured on an individual run basis and position discrimination was calculated from the full cross correlation matrix; the resulting values are necessarily slightly different from those in Figure 5, but show the same pattern of results. **b)** Correlation matrix for the five maps of BOLD response corresponding to the ignored stimulus positions, measured within the pulvinar. The correlation at each of the four stimulus separations was significantly greater than the bootstrapped null distribution (all $p < 0.001$, evaluated as the proportion of the bootstrapped null distribution that was greater than the actual correlation between ignored maps; bootstrapping based on 80 correlation values per subject), but the correlations did not show a systematic progression (slope) as a function of stimulus separation (Fisher $z = 0.18$; $p = 0.85$; nonparametric bootstrap test based on 50 points). **c)** Correlation matrices for the attended and ignored activity maps within bilateral V1. For both the attended and ignored stimulus positions, there was a systematic decrease in correlation between BOLD response maps with increasing stimulus separation (Fisher $z = -1.55$ for attended; $p < 0.001$; Fisher $z = -1.23$ for ignored; $p < 0.001$; nonparametric bootstrap tests based on 50 points per test).

Supplementary Methods:

Experimental design & analysis

In both experiments, we presented the stimuli in a blocked design. We recently explored the power of event-related vs. blocked designs for the correlation-based position discrimination analysis[20], and found that while both yielded qualitatively similar results, the blocked design was better suited to generating high SNR maps for the correlation analysis.

An important aspect of the design for both experiments is the simultaneous presence of attended and ignored stimuli, coupled with the up/down attention manipulation. Not only does the simultaneous presentation of targets and distractors induce competition that may be critical to engaging the pulvinar[3-6], it also allowed us to measure the precision of position coding for attended and ignored stimuli within the same set of voxels, and within the same functional runs. We conducted our analysis within regions of interest encompassing the entire pulvinar, and did not preselect the most active voxels as is sometimes done in classification analyses. Accordingly, the data that entered into the analyses for the attended and ignored locations was identical. Signal-to-noise ratio, vascular distribution, motion artifacts, and any other scanning artifacts were identical for the analyses of the attended and ignored stimuli, and cannot explain systematic differences between the precision of information that we measured in the two dimensions.

The up/down manipulation of attention (rather than left/right) was critical to the experimental design. Since the human pulvinar responds to stimulation in the contralateral visual field[12,14], by placing both a target and a distractor within the

same visual hemifield, we were able to test for encoding of attended and ignored stimuli within the *same* hemisphere, in the *same* functional runs. A left/right attention manipulation would confound the attended location with visual field; hence any comparison of the encoding of attended and ignored stimuli within the same hemisphere would have had to be done across runs, for stimuli that were not presented simultaneously. Our design excludes the possibility that any differences across runs or hemispheres contributed to the attentional gating we measured in the pulvinar.

Position discrimination analysis

In the position discrimination analysis, we used the central (9 deg. eccentricity) stimulus condition as a baseline for several reasons. First, doing so allowed us to fit a linear model to the resulting correlation plots. Our previous modeling has shown that the correlation between two patterns of BOLD response falls off nonlinearly with increasing separation between the corresponding stimuli[20]. By using the central map as a baseline, we use each of the four non-baseline maps exactly once, and reduce the risk of a nonlinearity or floor effect in the correlations washing out our ability to measure position discrimination. Second, using a single baseline increases the independence of correlations in the regression and reduces autocorrelations in the residuals (which would otherwise be a violation of the linear regression). Finally, and most importantly, the attentional modulation (gating) interaction (Fig. 5) still holds even when all position separations are analyzed (p = 0.004; bootstrap test for $z_{attended} > z_{ignored}$; n = 50 points for attended data and 50 points for ignored data). The results are

therefore not contingent on the use of a baseline condition. Using the central condition as a baseline, however, is a more conservative approach for the reasons described above and is therefore adopted in the analysis in Fig. 5.

A potential influence on our position discrimination measurements is retinotopically-specific adaptation—perhaps responses to the central stimulus position were reduced due to repeated stimulation at that location. In a control analysis, to compare the amplitude of the raw BOLD response across the five stimulus positions, we conducted ANOVAs for the five attended and five ignored conditions separately. The input to the ANOVAs was the average BOLD response within the pulvinar (t values), for each of the five stimulus positions, across 8 runs, for two hemispheres per subject, totaling 400 measurements. Each ANOVA used a mixed-effects model with fixed effects of hemisphere and stimulus position, and a random effect of subject, as well as the three two-way interactions among these factors. We found no adaptation effects in the raw BOLD response in the pulvinar: the BOLD response did not differ significantly across the five stimulus positions within the maps for the attended or ignored stimuli (main effect of stimulus position for attended stimuli: $F_{4,395} = 1.24$, $p = 0.29$; for ignored stimuli: $F_{4,395} = 0.65$, $p = 0.63$; based on five stimulus positions and 400 total measurements; no significant interactions with hemisphere or subject).

To establish a bootstrapped null distribution against which to compare the strength of the correlations we measured in the pulvinar and V1 (Supp. Fig. S1), we performed a label shuffling procedure. On each of 1000 iterations, we randomly shuffled the trial labels in the GLM design matrix for each run, then

recomputed the BOLD response map for each of the five stimulus positions. The shuffling included randomizing the fixation baseline labels, so as to remove all stimulus-related information from the GLM predictors. From the resulting five maps of BOLD response, we computed the pairwise correlations among the maps in the same way that we did for the position discrimination analysis in Experiment 1 (see Fig. 4), and recorded the mean correlation across all runs and all map pairs. Repeating this procedure 1000 times, we arrived at a distribution of correlation measurements that reflect a baseline level of correlation among response maps in our data set, independent of any stimulus-related information. This baseline level of correlation represents a floor for the correlations we might expect to find in the position discrimination analysis, and may result from factors such as partial voluming or local vascular structure.

Orientation classification analysis

The training stage of the classification analysis included a feature selection step that estimated an appropriate variance threshold to apply to the data before classifying (voxels below this threshold were excluded when classifying); the same variance threshold was then applied when testing the performance of the model. This training process, including variance threshold estimation, was performed completely independently of the test data, and introduced no information about the test data as confirmed by the label shuffling procedure described below.

Theoretical chance for the 2-class classification in Experiment 2 is 50%. To determine an empirical chance distribution for significance testing, we repeated the classification procedure 1000 times, each time shuffling the label assignments of the blocks (after running a GLM to recover the block-by-block beta weights, but before performing the classification step). The resulting distributions (measured separately for each subject and each ROI) reflected the range of classifier performance to be expected by chance. To test the significance of the classification accuracy for an ROI, we compared the actual classification accuracy with this permuted bootstrapped chance distribution, computing the proportion of the chance distribution that was larger than the actual classification accuracy.

Eye tracking

For three subjects, we monitored eye position during scanning. Eye position was recorded at 60 Hz using an ASL Eye-Trac 6 series long-range eye tracker and EyeTrac6000 software. If subjects' gaze position was selectively correlated with one of the dimensions of stimulus position (e.g., correlated with the attended positions but not the ignored ones, or vice versa), then eye movements could be a potential contributing factor to the differential encoding of attended vs. ignored positions that we measured in the pulvinar. To test for such a correlation, we binned the gaze position measurements by the stimulus that was present during the measurement, either at the attended or ignored location, and found no correlation between gaze position and either of the position dimensions (Fig. 2). Still, the possibility exists that there were some small but systematic eye

movements that we were unable to detect with our tracking. However, the data from V1 (Fig. 5c) rules out the possibility that any such eye movements could be responsible for our results in the pulvinar: V1 showed robust encoding of the ignored stimulus positions at almost the same degree of precision as for the attended stimuli. If the difference between the encoding of attended and ignored stimuli in the pulvinar was due to any kind of systematic eye movements, the same would also apply in V1. The fact that V1 and the pulvinar show a dissociation in their encoding of the ignored stimulus positions rules out the possibility that any kind of eye movements, voluntary or involuntary, could explain our results in the pulvinar.