Brainwide hemodynamics predict neural rhythms across sleep and wakefulness in humans

Leandro P. L. Jacob*[1], Sydney M. Bailes[1,3], Stephanie D. Williams[1,3], Carsen Stringer[4], Laura D. Lewis*[1,2,5]

[1] Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA

[2] Institute for Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, MA, USA

[3] Boston University, Boston, MA, USA

[4] HHMI Janelia Research Campus, Ashburn, VA, USA

[5] Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Boston MA USA

* Corresponding authors: ljacob@mit.edu – ldlewis@mit.edu

Keywords: Sleep, Arousal, EEG, fMRI, Machine Learning

Abstract:

The brain exhibits rich oscillatory dynamics that play critical roles in vigilance and cognition, such as the neural rhythms that define sleep. These rhythms continuously fluctuate, signaling major changes in vigilance, but the brainwide dynamics underlying these oscillations are unknown. Using simultaneous EEG and fast fMRI in humans drifting between sleep and wakefulness, we developed a machine learning approach to investigate which brainwide fMRI networks predict alpha (8-12 Hz) and delta (1-4 Hz) fluctuations. We predicted moment-to-moment EEG power variations from fMRI activity in held-out subjects, and found that information about alpha rhythms was highly separable in two networks linked to arousal and visual systems. Conversely, delta rhythms were diffusely represented on a large spatial scale across the cortex. These results identify the large-scale network patterns that underlie alpha and delta rhythms, and establish a novel framework for investigating multimodal, brainwide dynamics.

INTRODUCTION

As the brain navigates tasks and states, neurons frequently fire in synchrony, generating rich oscillatory dynamics with varied temporal and spatial distributions. These neural rhythms have been studied with scalp EEG for almost a century, and have formed a core aspect of modern neuroscience due to their established links to cognitive computations and brain states. Neural rhythms are particularly fundamental to the study of vigilance states and sleep, which are defined by the appearance of distinct oscillatory EEG patterns with well-studied physiological outcomes [1], including several forms of memory [2–10] and attention [11,12], along with basic physiological processes such as brain waste clearance [13]. Extensive work has been dedicated to identifying the oscillatory generators behind these electrophysiological rhythms. However, the waxing and waning of neural rhythms can also reflect global shifts in brain state, beyond the specific subset of neurons oscillating at a given frequency. While EEG is widely used to measure arousal and attentional states, the brainwide activity patterns that underlie these EEG rhythms are not established.

Two neural rhythms with key relevance for the study of arousal and cognition are delta (1-4 Hz) and alpha (8-12 Hz). Alpha (8-12 Hz; Fig. 1a) is implicated in an array of cognitive processes [12,14,15], and maximally detected in occipital electrodes during eyes-closed wakeful rest. Delta (1-4 Hz; Fig. 1a) is most prominent during non-rapid eye movement (NREM) sleep and strongly linked to its cognitive benefits [4,5,9,16]. Prior work has sought to identify the focal generators of these oscillations [17–25]. However, the brainwide dynamics that they represent—especially the dynamics that modulate oscillatory generators, but do not oscillate at the same frequency—are less clear. Specifically, when alpha or delta rhythms fluctuate as attention and vigilance shift, what brainwide activity does this represent, and which specific networks are affected?

Identifying the brainwide dynamics that underlie variations in neural rhythms is technically difficult to investigate: scalp EEG suffers from low spatial resolution and is unable to resolve deep brain activity, while invasive methods such as electrocorticography have limited spatial coverage. Simultaneously acquiring EEG and functional magnetic resonance imaging (fMRI) can provide high spatial resolution and the ability to image deeper brain regions. Analyzing EEG-fMRI data, however, poses a significant challenge: the BOLD (blood-oxygenation-level-dependent) signal measured by fMRI is an indirect measure of brain activity, and while it can be modeled using a hemodynamic response function (HRF), the BOLD signal's true relationship to the underlying neural signal is complex. The coupling between BOLD and neural activity varies across brain areas [26], and the HRF fails to account for many temporal properties of fMRI signals [27–29]. Alternative approaches that fit a more flexible HRF are susceptible to noise and may overfit the data. To avoid overfitting, cross-validation is essential, as even traditional correlation approaches using linear regression may overfit and not be generalizable to the greater population [30]. Finally, standard methods primarily conduct univariate analyses that consider each region or voxel separately, which can identify individual regions correlated with neural rhythms, but leave open the question of how joint activity across brain areas is linked to neural dynamics. Predictive methods have become a powerful approach in fMRI [31,32] as they overcome these limitations: they can learn network-level information across multiple brain regions with flexible data-driven functions, and they inherently involve cross-validation to ensure results are generalizable.

These approaches have identified relationships between neural activity and fMRI signals [33], and can predict a vigilance index from fMRI [34,35]. However, whether specific EEG rhythms can be predicted from fMRI signals is not yet established.

We designed a predictive framework to identify how brainwide dynamics are linked to distinct EEG rhythms, using machine learning approaches to predict alpha and delta power from fMRI data. We used two independent datasets of simultaneous EEG and accelerated fMRI (TR<400 ms), exploiting the high temporal information of our fMRI data to model fluctuations in neural rhythms as subjects spontaneously drifted in and out of sleep. Critically, our goal was not to localize focal oscillatory sources, but rather to identify the brainwide activity patterns that signal variations in EEG rhythms, as scalp oscillations can reflect brain dynamics that need not oscillate at the same frequency. We first show that we can successfully predict neural rhythms from hemodynamics in out-of-sample subjects. We find that distinct and replicable fMRI patterns predict each rhythm, and that predictive information is present even at the single-region level. Finally, we develop an analysis to probe the joint fMRI information across brain regions, and use it to identify distinct network patterns underlying each neural rhythm. This approach identifies how fMRI information predictive of neural rhythms is spatially distributed across the brain, separating networks of regions with shared (vs. distinct) information about EEG dynamics. We found a large, diffuse network of primarily cortical regions that predicted variations in delta power, whereas alpha power was predicted by two separable networks primarily representing the distinct dynamics of the visual system and subcortical arousal control circuits.

RESULTS

EEG alpha and delta power were predicted from brainwide fMRI data in held-out subjects

We first investigated whether it is possible to predict the continuous fluctuations of neural rhythms from simultaneously acquired fast fMRI data. Prior work successfully predicted EEG-based vigilance metrics (a ratio between two frequency bands) from simultaneous fMRI data and achieved a mean Pearson's linear correlation value of r=0.31 between predictions and targets when training on the same subjects [35]. However, predicting distinct rhythms is a more challenging task, as fluctuations in individual frequency bands are noisier and more complex than global variations in vigilance metrics, which collapse multiple distinct EEG properties into a single simplified metric. Additionally, our goal of generalizing to out-of-sample subjects could make predicting individual frequency bands particularly difficult, as EEG power magnitudes exhibit strong individual differences. Thus, based on prior work and the variability of individual EEG, we expected a performance ceiling in the range of r=0.3 (Pearson's linear correlation between held-out subject predictions and ground truth). We hypothesized that by designing a flexible machine learning approach and training with the high temporal resolution of fast fMRI, we could predict time-varying EEG oscillations.

We tested our models on held-out subjects (not seen during training) to ensure overfitting did not inflate our performance estimates and to assess the generalizability of the model across individuals. We employed two separate datasets collected at different sites, to further confirm reproducibility (dataset A: n=21, TR=0.378s; dataset B: n=11, TR=0.367s). For alpha analyses, spectrograms were first inspected to identify whether subjects had clear alpha rhythms, since some individuals do not exhibit typical eyes-closed alpha [36,37] (dataset A: n=18; dataset B: n=6). We trained linear machine learning models to predict EEG alpha and delta power, calculated in sliding 5-second windows from occipital electrodes (which are less susceptible to MR-induced noise), from simultaneously collected fMRI data, in subjects drifting in and out of sleep (Fig. 1b-c). The fMRI data was parcellated in 84 regions, consisting of 31 bilateral cortical regions, 7 bilateral subcortical regions, and 8 non-gray matter regions (white matter and ventricles).

The model was trained on sliding windows (length of 60 TRs, ~20s) of the parcellated fMRI data, with each segment yielding a single-point prediction of normalized EEG power (at the temporal center of the window). We centered the window of fMRI predictors on the EEG point being predicted to allow the model to flexibly learn whichever fMRI temporal patterns were most relevant, rather than assuming a specific lag between EEG and fMRI. The model was thus trained on 5,040 predictors (84 regions with 60 time points per region). Training was carried out separately for each dataset,

iteratively on all subjects but one, with testing then performed on the held-out subject. The performance was compared to a shuffled control, breaking the true relationship between the modalities.

This whole-brain model successfully predicted EEG power in out-of-sample subjects for both alpha ('brainwide' vs 'control', Fig. 1d, $t(23)=4.7$, $p<0.001$) and delta ('brainwide' vs 'control', Fig. 1f, $t(31)=14.8$, $p<0.001$). Delta predictions surpassed the expected performance ceiling, displaying a mean correlation between predictions on held-out subjects and ground truth of $r=0.41$, and alpha was predicted at $r=0.30$. The models thus robustly predicted fluctuations of neural rhythms in held-out subjects, demonstrating the ability to learn dynamics shared across individuals. Representative predictions demonstrate that the model successfully captured short- and long-timescale fluctuations in alpha and delta power (Fig. 1e, 1g).

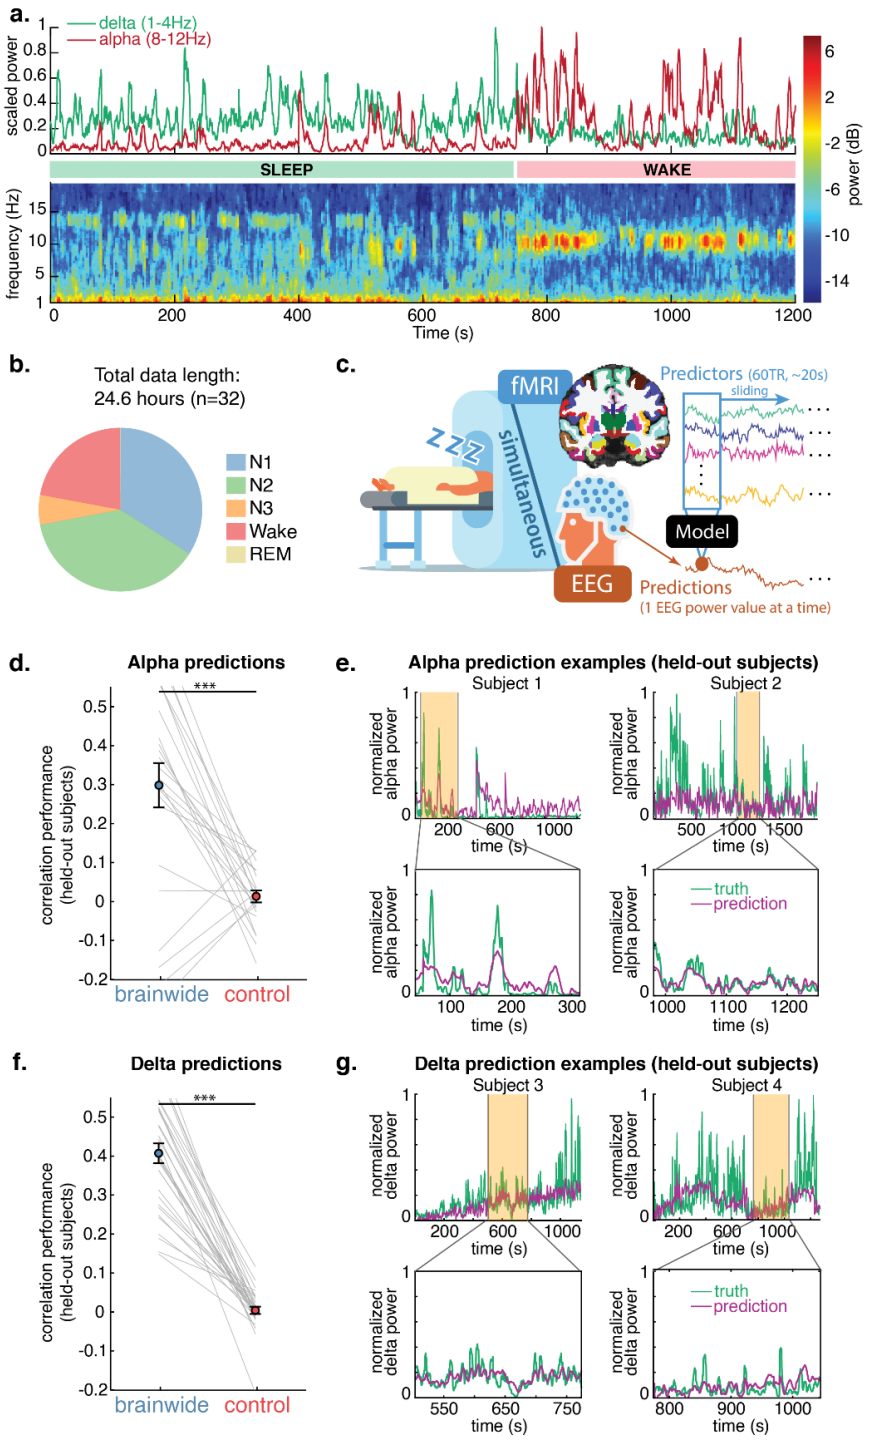

Fig. 1: Machine learning can predict dynamic fluctuations in neural rhythms using brainwide fast fMRI timeseries. a. *Example of occipital alpha and delta EEG power fluctuating across sleep and wakefulness in one subject.* ***b.*** *Sleep stage distribution show subjects drifted between wake and sleep (primarily N1 and N2).* ***c.*** *EEG alpha and delta power were separately predicted from simultaneous fMRI data, collected while subjects rested inside the scanner with their eyes closed, drifting in and out of sleep. fMRI data was parcellated into 84 cortical, subcortical, and non-gray matter regions. Each point in the EEG power series, interpolated to match TR times, was predicted by 60 TRs of parcellated fMRI data.* ***d.*** *Correlation between alpha predictions (on held-out subjects) and ground truth, demonstrating that predictions using brainwide fMRI data are significantly better than the control condition, in which the fMRI data was shuffled within each subject run. *** p<0.001, two-tailed paired t-test. Gray lines indicate held-out subject prediction performance. Error bars are SEM; N=24.* ***e.*** *Examples of alpha prediction in individual subjects, showing tracking of short- and long-timescale alpha power fluctuations.* ***f.*** *Delta prediction performance, across the same conditions as panel d, shows significantly above-chance delta prediction. N=32.* ***g.*** *Examples of delta prediction in individual subjects.*

Cortical, subcortical, and non-neural fMRI signals differentially predict alpha and delta rhythms

Given that we were able to predict neural rhythms on out-of-sample subjects, we next explored how we could use this approach to identify which regions carry information about alpha and delta fluctuations. Importantly, both neuronal and non-neuronal signals (such as blood vessels and cerebrospinal fluid) carry information about general fluctuations in arousal state [13]. We therefore first examined whether this predictive information was specific to gray matter regions in cortex or subcortex, or whether it was also carried in non-neuronal regions such as the white matter and the ventricles. We trained new models under three conditions: cortex only, subcortex only, and non-gray matter regions only, with testing again performed on held-out subjects (Fig. 2a). These conditions were then compared with the 'brainwide' and 'control' (fMRI data shuffled within each run) conditions described above.

Alpha power (Fig. 2b, Extended Data 1) was successfully predicted above chance by cortical (r=0.27, p=0.002) and subcortical (r=0.31, p<0.001) fMRI signals, with prediction performance achieving a similar level as the brainwide condition (r=0.30). The finding that subcortical regions alone could predict alpha rhythms on a held-out subject with such high accuracy was surprising, as these regions have a lower signal-to-noise ratio. Finally, alpha power could not be significantly predicted by the non-gray matter regions on their own, indicating that alpha-predictive fMRI activity relied on the neural information present in the BOLD signal.

Delta power (Fig. 2c, Extended Data 2), on the other hand, could not only be predicted by neural sources, but was also significantly predicted (with worse performance) by the non-gray matter regions on their own (r=0.24, p<0.001), likely owing to established coupling between delta power, cerebrospinal fluid flow (in the ventricles), and global blood flow [13]. In other words, the delta-predictive fMRI signals possessed a significant non-neural component; however, the brainwide delta-predictive fMRI information also relied on neural information, as the non-gray matter condition generated worse predictions than the brainwide condition (r=0.41, p<0.001) and the cortical condition (r=0.37, p=0.001). The cortical condition achieved similar predictive performance as the brainwide condition, but the subcortical condition performed significantly worse (r=0.29, p=0.01), indicating that delta predictions did not benefit as strongly from subcortical information as alpha did.

Given that our study employed two distinct datasets, collected at different sites, we also analyzed each dataset independently (Extended Data 3). The distinct predictive patterns for alpha and delta were strikingly consistent across the datasets.

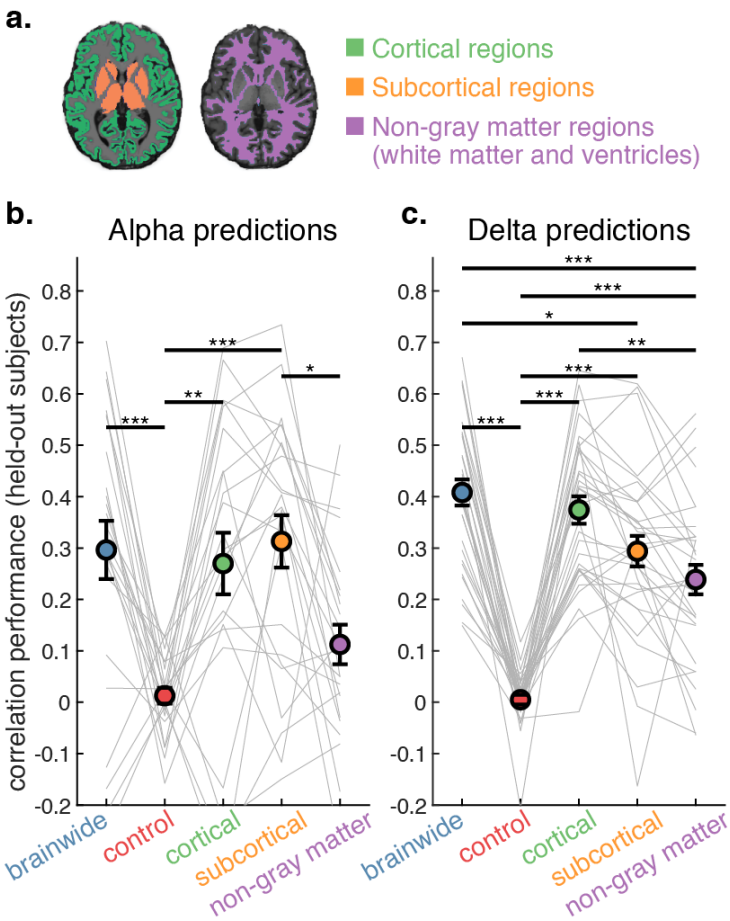

Fig. 2: Alpha and delta rhythms show distinct relationships to cortical, subcortical, and non-neural fMRI signals. ***a.*** *Models were trained and tested under three additional conditions: cortical regions only, subcortical regions only, and non-gray matter regions only (ventricles and global white matter). Performance was compared to models trained using all regions, and under a control condition using all regions but with shuffled fMRI data, breaking the true relationship between EEG and fMRI.* ***b.*** *Alpha power can be predicted by cortical and subcortical regions, but not by non-gray matter regions. Circles show mean correlation between alpha power ground truth and model predictions (on held-out subjects). * p<0.05; ** p<0.01; *** p<0.001; Tukey's HSD. Error bars are SEM. Gray lines indicate individual held-out subject prediction performance.* ***b.*** *Delta power can not only be predicted by cortical and subcortical regions, but also by non-gray matter regions on their own.*

Individual fMRI regions can predict neural rhythms, with distinct spatial patterns for alpha and delta

To next identify finer-scale information beyond the coarse anatomical divisions of cortex and subcortex, we tested the information contained within each region in isolation by training and testing the model on a single bilateral gray matter region at a time. This prediction approach allows for the modeling of flexible temporal functions shared across subjects linking each neural rhythm to each fMRI area. We generated spatial maps of individual brain regions' predictive performances for alpha and delta, to visualize the large-scale pattern of information content across the brain (Fig. 3). For visualization, the delta results combined the performance on held-out subjects for both datasets (n=32; model training was carried out independently in each dataset due to different imaging parameters). Alpha spatial maps included only Dataset A (n=18) as only 6 subjects in Dataset B displayed alpha rhythms, too few to assess predictions when using only a single region. These maps demonstrated that single brain regions could be sufficient to achieve successful prediction of neural rhythms, but with poorer performance than with a larger network. A control analysis (Extended Data 4) confirmed the significance of the global predictive maps by comparing its patterns to shuffled data.

This individual region analysis re-confirmed the predictive power of subcortical regions for alpha rhythms (Fig. 3a, Extended Data 5), revealing that alpha could be significantly predicted using only the thalamus (r=0.27, p<0.001), the dorsal striatum (separately as the putamen, r=0.23, p=0.002, and caudate, r=0.21, p=0.008), the hippocampus (r=0.20, p=0.002), or the pallidum (r=0.18, p=0.01). This result highlighted that neural rhythm prediction can be successful even in individual deep brain regions with lower signal-to-noise ratio. Consistent with the established coupling between alpha oscillations and fMRI activity within the visual system [17–22], cortical areas capable of individually predicting alpha power primarily comprised regions that process visual information (see Extended Data 5 for all correlation values and p-values), including early visual cortex (encompassing the cuneus, lingual, and pericalcarine), higher-order visual regions (fusiform, inferior temporal, parahippocampal, superior parietal, precuneus), and multisensory regions (superior temporal and its bank, and transverse temporal [38]). Alpha rhythms were also successfully predicted by the anterior cingulate cortex (separately as caudal and rostral), a region linked to arousal and often correlated with thalamic activity [39]. Prior EEG-fMRI studies identified correlations between alpha rhythms and the thalamus and occipital cortex [17–22]; our results demonstrated that alpha fluctuations also reflect changes in neural activity in additional arousal-coupled cortical and subcortical regions, and to regions downstream of occipital cortex that process visual information (including the hippocampus), highlighting this predictive approach's ability to learn more complex functions relating EEG and fMRI signals.

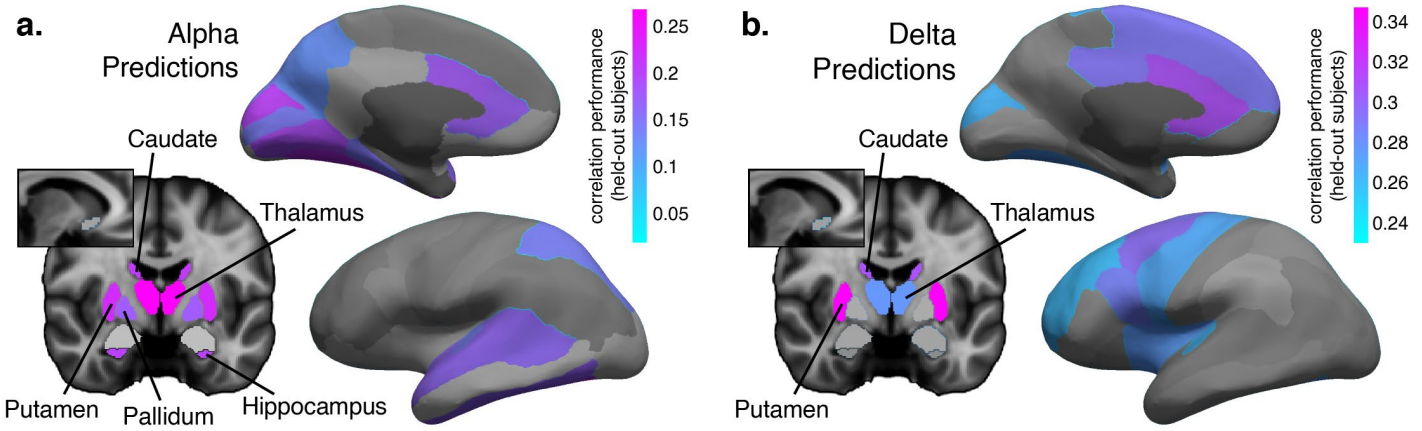

Fig. 3: Prediction performance maps show which regions individually carry significant information about alpha and delta rhythms. a*. Prediction performance (mean correlation between alpha predictions on held-out subjects and ground truth) when model was trained on single bilateral regions to predict alpha power. Regions shown in color yielded alpha predictions that were significantly better than control (non-significant regions are shown in grayscale). See Extended Data 5 for a list of correlation performance values.* ***b.*** *Prediction performance for delta. Model was trained on single bilateral regions combined with non-gray matter areas to predict delta power, in order to account for the non-neural delta-predictive components (see Fig. 2). Regions shown in color yielded predictions that were significantly better than non-gray matter areas alone. See Extended Data 6 for a list of correlation performance values.*

For the delta analysis (Fig. 3b), in order to identify delta-predictive information that is specifically derived from neural sources, we accounted for the fact that delta power could be predicted from non-neural regions alone (Fig. 2), by using this non-neuronal delta model as a control condition. For each gray matter region, we trained a model that included both the region in question along with the non-neural regions, and then compared its performance to the model that used only the non-neural regions. This approach allowed us to identify gray matter regions that contain unique delta-predictive information not present in the non-neuronal components of the fMRI signal, while avoiding introducing errors from regressing out physiological confounds [40–42]. Delta fluctuations were predicted most strongly by the putamen (r=0.35, compared to the non-neural model, r=0.24; p<0.001), again demonstrating the predictive power of subcortical regions despite their lower signal-to-noise ratio. The caudate (r=0.31, p=0.001) and thalamus (r=0.28, p=0.04) also contained delta-predictive information significantly above the non-neural control. The cortical delta-predictive information was distributed distinctly from alpha, overwhelmingly favoring the anterior part of the brain (an initiating site for delta

oscillations [43]), including the anterior cingulate regions, which have previously been implicated in delta-band wave propagation [43].

Two networks predict alpha, while delta information is diffuse

Our individual-region approach demonstrated that predictive information is present even on a single-region level, with distinct predictive maps for alpha and delta rhythms. However, the rise and fall of neural oscillations likely reflects changes in brainwide network dynamics spanning both the cortex and subcortex; thus, single-regions predictive maps could miss integrated networks that carry joint information about EEG oscillations. We therefore developed an analysis to identify predictive information shared across brain regions (Fig. 4a). We first assessed the predictive performance of each pairwise combination of fMRI regions in relation to the performance of its individual components, in order to evaluate whether regions contained distinct and complementary information (i.e., greater performance for the pair) or similar and redundant information (i.e., same or worse performance for the pair). In other words, we trained and tested the model on each pair of regions, and investigated the performance benefits (or losses) of each pair of regions (pair performance is shown in Extended Data 7), as compared to the individual region performance (Fig. 3). As before, non-gray matter regions were also used when generating delta predictions, to identify neural-specific fMRI information beyond that present in the white matter or ventricles, examining the same data as in the prior analysis. Next, we performed k-means clustering on these results, using the gap statistic to evaluate the optimal clustering solution (including the null model of no suitable clustering) [44]. By clustering the pairwise performance benefits, we could identify the presence of different networks of information: groups of regions that share redundant information among themselves (resulting in performance losses or no benefits when paired together), but which generate performance benefits when paired with regions from other predictive networks.

We found that alpha pairwise performance benefits were optimally clustered in three groups (Extended Data 8). Two of these groups (Fig. 4b, blue and red clusters), when paired together, displayed substantial performance benefits (r=0.058, *95% CI* [0.054, 0.062]). Meanwhile, the third group (Fig. 4b, green cluster) included only regions that could not significantly predict alpha power (see Fig. 3a), and generated virtually no performance benefits when paired with either of the other groups (r=0.007, *95% CI* [0.002, 0.012]; r=0.008, *95% CI* [0.005, 0.011]). None of the clusters generated benefits when paired with itself (r=-0.007, *95% CI* [-0.01, -0.003]; r=-0.004, *95% CI* [-0.007, -0.002]; r=-0.011, *95% CI* [-0.015, -0.007]). These results indicated that the blue and red clusters represented two distinct networks of with different alpha-predictive information, while the green cluster represented regions which do not contain appreciable alpha-predictive information. A control analysis (Extended Data 9) found that correlation between fMRI regions (as in typical functional connectivity analyses) could not explain the clustering behavior, confirming that this analysis identified network structures distinct from traditional fMRI analyses.

The smaller cluster of alpha-predictive information (Fig. 4b, blue) included the thalamus, the dorsal striatum (putamen and caudate), the pallidum, the anterior (separately as caudal and rostral) and posterior cingulate cortices, and the transverse temporal cortex. These results suggest that one component of alpha predictions reflected activity variance in this network of regions linked to arousal state modulation [39,45]. The larger cluster of alpha-predictive information (Fig. 4b, red), included visual cortex (both the ventral and dorsal visual streams), prefrontal and temporal regions associated with higher-order visual cognition, and three subcortical regions that are responsive to contents of visual information (the accumbens, amygdala and hippocampus), suggesting a distinct category of alpha-predictive information present in a network of regions involved in visual processing.

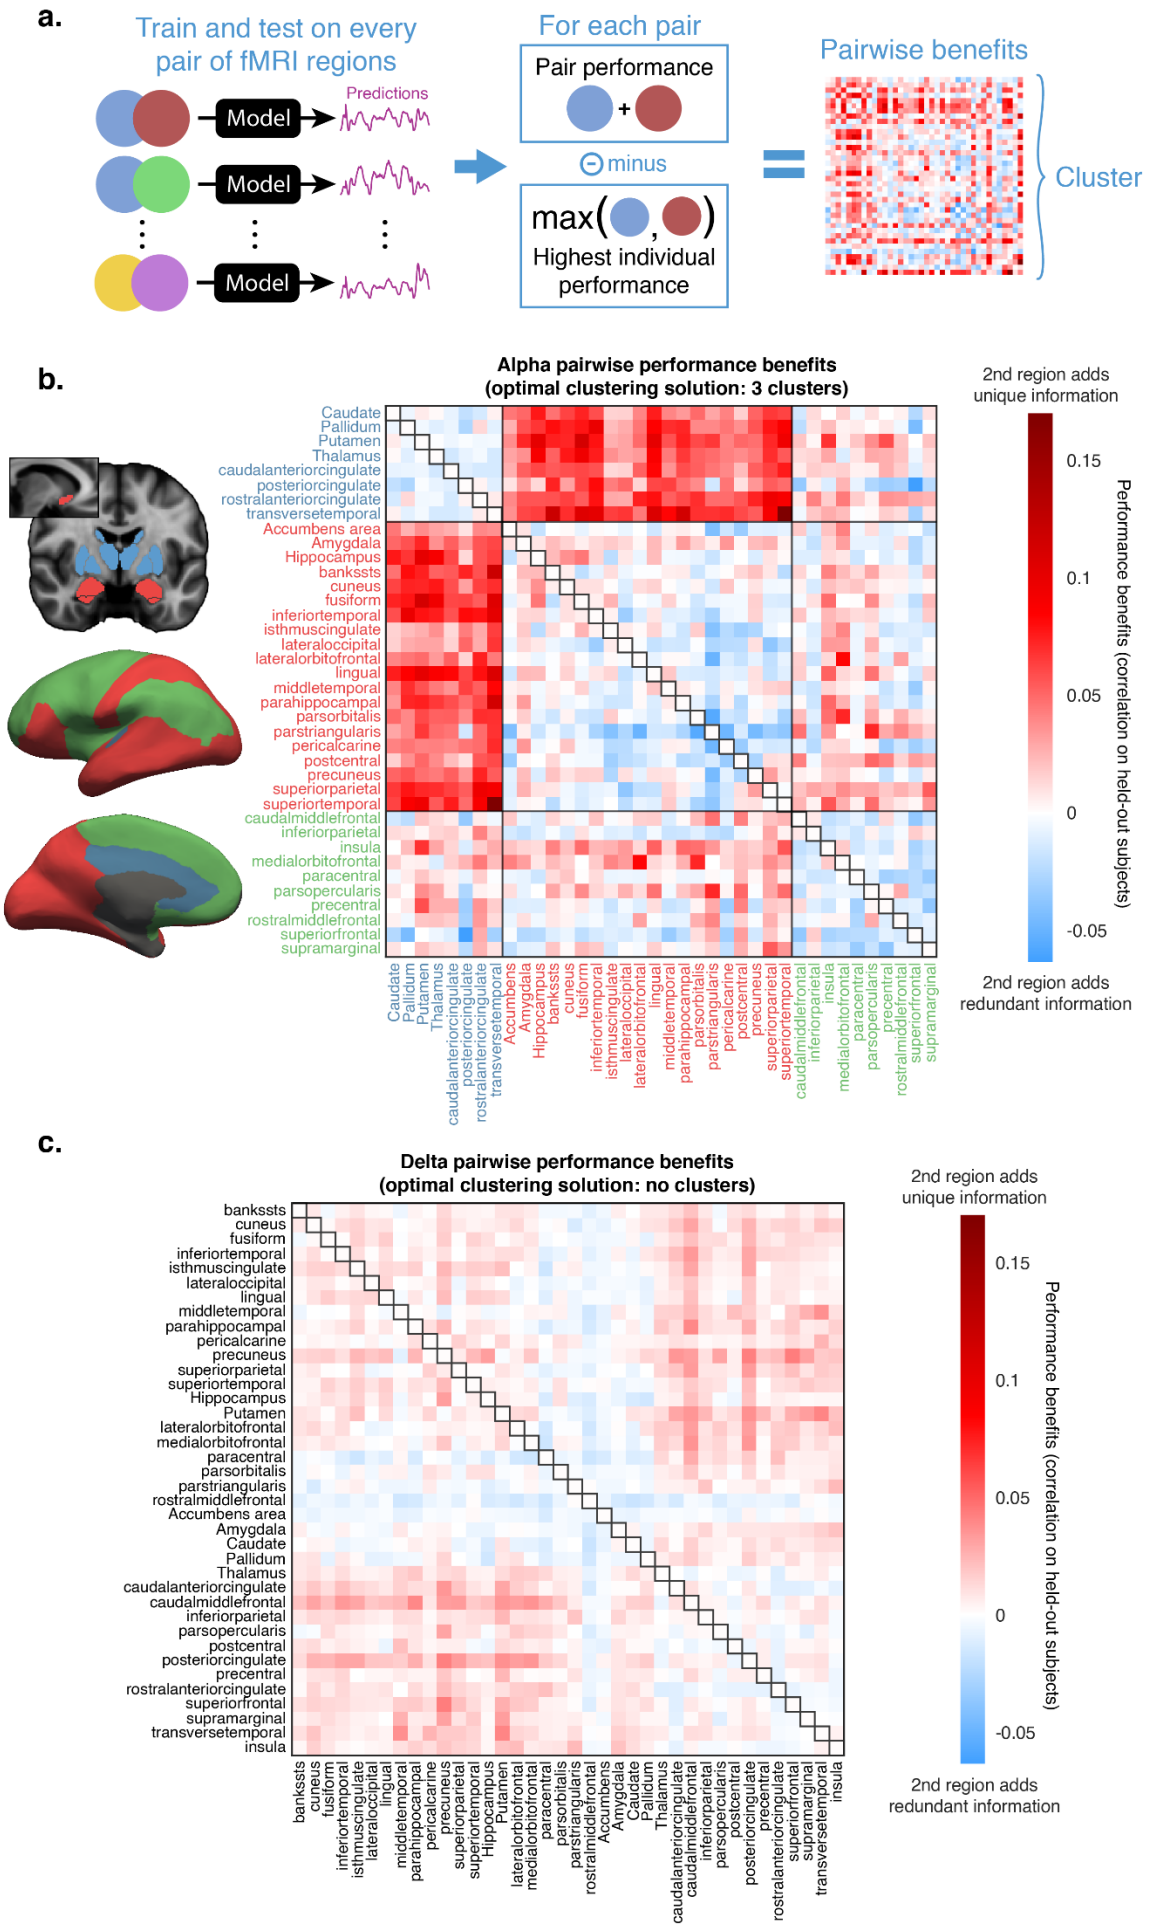
a.
Train and test on every pair of fMRI regions
Model
Predictions
For each pair
Pair performance
minus
max(,)
Highest individual performance
Pairwise benefits
Cluster
b.
Alpha pairwise performance benefits
(optimal clustering solution: 3 clusters)
2nd region adds unique information
2nd region adds redundant information
Performance benefits (correlation on held-out subjects)
0.15
0.1
0.05
0
-0.05
Caudate
Pallidum
Putamen
Thalamus
caudalanteriorcingulate
posteriorcingulate
rostralanteriorcingulate
transversetemporal
Accumbens area
Amygdala
Hippocampus
bankssts
cuneus
fusiform
inferiortemporal
isthmuscingulate
lateraloccipital
lateralorbitofrontal
lingual
middletemporal
parahippocampal
parsorbitalis
parstriangularis
pericalcarine
postcentral
precuneus
superiorparietal
superiortemporal
caudalmiddlefrontal
inferiorparietal
insula
medialorbitofrontal
paracentral
parsopercularis
precentral
rostralmiddlefrontal
superiorfrontal
supramarginal
c.
Delta pairwise performance benefits
(optimal clustering solution: no clusters)
2nd region adds unique information
2nd region adds redundant information
Performance benefits (correlation on held-out subjects)
0.15
0.1
0.05
0
-0.05
bankssts
cuneus
fusiform
inferiortemporal
isthmuscingulate
lateraloccipital
lingual
middletemporal
parahippocampal
pericalcarine
precuneus
superiorparietal
superiortemporal
Hippocampus
Putamen
lateralorbitofrontal
medialorbitofrontal
paracentral
parsorbitalis
parstriangularis
rostralmiddlefrontal
Accumbens area
Amygdala
Caudate
Pallidum
Thalamus
caudalanteriorcingulate
caudalmiddlefrontal
inferiorparietal
parsopercularis
postcentral
posteriorcingulate
precentral
rostralanteriorcingulate
superiorfrontal
supramarginal
transversetemporal
insula

Fig. 4: Clustering analysis of pairwise performance reveals two distinct networks of alpha-predictive fMRI information, while delta-predictive information is not separable in clusters. a. *Models were trained on every possible pair of non-gray matter regions, and correlation performance was calculated on held-out subjects. Pairwise performance benefits were calculated as the performance gain (or loss) of a pair in relation to the maximum performance of its individual components; see Fig. 3 for individual performance values. Performance benefits were then clustered and evaluated with the gap statistic (Extended Data 8) to determine optimal clustering solution.* ***b.*** *Identified clusters and all pairwise performance benefits for alpha-predictive information. When regions from the blue and red cluster were paired together, they consistently obtained performance benefits not found in any other cluster match.* ***c.*** *The clustering analysis revealed that delta did not possess a suitable clustering solution, as evidenced by no systematic pairwise performance benefits following a forced 3-cluster fitting.*

In contrast, the gap statistic demonstrated that the delta results did not present a suitable clustering solution (Extended Data 8). To illustrate the distinct network structures that represent alpha and delta, we forced a three-cluster fitting on the pairwise delta predictions (Fig. 4c), showing that delta did not display the clustering behavior present in the alpha analysis (Fig. 4b). Delta-predictive information, therefore, was not expressed as separable networks.

To obtain further insights on the network structure behind these rhythms, and particularly to better understand why delta did not possess a suitable clustering solution, we conducted an additional analysis to assess the size of fMRI networks required to predict each EEG rhythm. In other words, how many distinct fMRI regions are needed for optimal predictions? If many regions are required, this implies that predictive information is widely distributed across the brain, with low redundancy between regions. On the other hand, if only a small subset of regions is required for optimal predictions, this suggests that predictive information is largely redundant across regions. To identify this number of optimal regions, we designed an iterative procedure to remove redundant information: in this procedure, low-weight regions are iteratively removed from the model (starting with all 84 regions) while performance changes are tracked.

We found that removing low-weight regions enabled us to achieve improved predictions of both alpha and delta power in relation to using brainwide data (Extended Data 10). These performance increases peaked at a larger number of regions for delta than for alpha (37 regions, vs. 18 regions), and subsequent reductions in region numbers caused performance to drop, eventually reverting the benefits of feature selection. This turning point happened only at very low numbers for alpha predictions (8 regions), indicating that alpha-predictive information has low dimensionality. This finding was consistent with alpha possessing a small number of separable networks of shared information. Meanwhile, for delta predictions, the model degraded at a much larger number of regions (17 regions) with a more pronounced performance loss afterwards, indicating that unique delta information is present in a large number of regions with low redundancy between regions. Finally, performance improvements at the optimal iteration were more pronounced for alpha (r=0.46 when using 18 regions, compared to brainwide r=0.30, $t(17)=3.32$, p=0.008) than delta (r=0.45, compared to brainwide r=0.41, $t(31)=4.30$, $p<0.001$), further pointing to alpha benefitting more strongly from feature selection due to the high redundancy of its representation. Taken together, these analyses revealed two highly separable networks of alpha-predictive information, representing the dissociable components of arousal and visual dynamics, with a large amount of information redundancy within each network. Meanwhile, delta information was highly diffuse, requiring a large array of brain regions for optimal predictions.

DISCUSSION

Delta and alpha rhythms signal changes in vigilance and play major roles in human cognition, yet the underlying network patterns they represent—beyond focal oscillatory generators—have been unclear. We used machine learning to predict continuous fluctuations in EEG power in the alpha and delta bands from brainwide fMRI data, and achieved robust performance in held-out subjects, with a mean correlation value of r=0.41 and r=0.30 (respectively for delta and alpha; Fig. 1) between predictions and ground truth using brainwide data, and r=0.45 and r=0.46 following feature selection (Extended Data 10). This prediction-based approach demonstrated reproducibility across datasets (Extended Data 3), and was able to uncover information about continuous neural rhythm variations in individual fMRI regions (Fig. 3). We discovered that distinct network patterns were coupled to each rhythm, with strikingly differentiated spatial scales: two

focused, separable networks predicted alpha rhythms from moment to moment, whereas a highly diffuse array of primarily cortical regions provided information about delta oscillations.

Alpha rhythms have been widely studied in the EEG-fMRI literature, with focus on identifying its specific generators by using mass-univariate approaches to correlate individual fMRI regions to alpha power variations. Our findings replicate prior univariate relationships between the thalamus/occipital cortex and alpha power [17–24]. We also identified novel relationships between alpha power and fMRI activity in several cortical and subcortical regions, by performing predictions without requiring a pre-defined HRF. Our results demonstrate that separable networks are coupled to alpha rhythms, suggesting they reflect two distinct mechanisms of occipital alpha expression: the dynamics within the visual system itself, and the modulation from the arousal system. The thalamus has often been viewed as the primary orchestrator of alpha rhythms [18–20,24], and our results indicate that this thalamic control appear to be part of a circuit involving other key arousal regions: the dorsal striatum, the pallidum (which receives projections from the dorsal striatum, and in turn projects to the thalamus [46]), and the cingulate cortex (which closely communicates with the thalamus [47,48]). Meanwhile, the separable alpha-predictive information identified in the visual system points to a distinct and dissociable mechanism of alpha expression, which may reflect that cortical neurons can intrinsically fire in the alpha range independent of thalamic input [49], and that alpha rhythms have been proposed to be involved in visual information flow from lower-level perceptual areas to higher-order cognitive regions [15,50].

We found a distinct network pattern linked to delta power, which indexes several sleep oscillatory dynamics, including slow oscillations and slow waves. While prior EEG-fMRI studies have identified fMRI activity locked to individual slow wave events [51–53], the delta power variations examined here represent distinct information, and past work did not find significant correlations between delta and fMRI activity [21]. Our work explains this null result by showing that delta predictions rely on combined information across a large number of fMRI regions, while also possessing a substantial non-neuronal component. These results can be contrasted with alpha, which is readily predicted by a small network, and cannot be significantly predicted by non-gray matter fMRI regions.

Delta predictions not only relied on a large network, but also did not form separable clusters. This suggests that subtle differences in fMRI activity, primarily across the cortex, came together to predict variations in delta rhythms, as highlighted by the coarse-grained analysis (Fig. 2c) identifying superior delta predictions when using the full conjunction of cortical regions. Thus, delta predictions represented an accumulation of diffuse activity primarily distributed throughout the cortex. What dynamics within this broad network, in which each region possesses unique information, could be predicting delta power? Prior work has shown that arousal fluctuations are associated with cortical infra-slow propagating fMRI activity [54], and that infra-slow fMRI is coupled to delta power [55]; it is therefore possible that this propagating infra-slow fMRI activity was core to the delta predictions. The varying phase across brain regions may explain why delta-predictive information required a large number of regions and was not separable in different networks. In addition, delta power itself exhibits complex propagating spatiotemporal dynamics [56], and can be both local and global in nature [5,57], suggesting that this diffuse fMRI predictive network may also represent a collection of diverse local slow wave dynamics. Future work could explore these questions by identifying traveling wave events in simultaneous EEG-fMRI and investigating whether distinct spatiotemporal features of infra-slow fMRI oscillations predict local or global delta power variations.

Past EEG-fMRI analyses of alpha and delta rhythms have primarily relied on univariate correlations, considering each brain region or voxel separately [17–21,23], yet fMRI activity is strongly characterized by network-level activity, with variations in multiple regions generating distinct spatiotemporal patterns [58,59]. Our approach allows us to not only identify which regions are individually coupled to neural rhythms, but to characterize whether the information across regions is unique or redundant, enabling us to identify networks. As an example, we found that both the transverse temporal and superior temporal cortices, regions that are most commonly associated with auditory perception, were individually predictive of alpha rhythm variations to a similar degree. However, we found that they were in separate networks, with the superior temporal cortex—known to also play a role in visual perception [38]—being in fact part of a network of visual alpha dynamics. Meanwhile, the transverse temporal cortex shared the same category of alpha-predictive information as an array of subcortical arousal-controlling regions, which may be explained by its activity reflecting thalamic-mediated state

modulation, with prior work finding that the transverse temporal cortex is affected by sleep spindle density [60] and poor sleep quality [61].

Predictive machine learning has supported neuroscientific insights in event-related classification analyses across a range of neuroimaging modalities [62–65], and prior work has shown that some localized fMRI activity can be predicted from the EEG signal [66,67]. Our method extended these predictive frameworks to achieve continuous, dynamic predictions of neural rhythms (which are much higher frequency than the fMRI temporal resolution) from brainwide fMRI data, in the absence of external stimuli and task conditions. This approach has several advantages, such as correcting for non-neuronal effects (by including the white matter and ventricles as predictors) without requiring regression that can introduce errors [40–42], ensuring that effects are generalizable across subjects through cross-validation, and allowing for the identification of multi-region relationships that represent network-scale information. fMRI studies often use hemodynamic functional connectivity (defined by correlation patterns) to identify network dynamics; while functional connectivity can reflect coherent neural activity, other aspects of electrophysiological networks are not captured in this metric, and ECoG work has shown that distinct neural frequency bands form spatially specific networks [68]. This property is well aligned with our results that alpha and delta information have distinct underlying networks that cannot be detected by correlation patterns between fMRI regions. Finally, this predictive approach can be generalized to other continuous crossmodal analyses in neuroimaging data, to identify how spontaneous brain dynamics are coupled in a broad range of questions. Future work could identify distributed spatial dynamics underlying additional brain rhythms, investigate how these networks of information may vary across cognitive tasks, or probe how these networks may shift in aging and disease. This predictive approach could also be adapted to investigate which specific temporal patterns in fMRI predict neural rhythms.

In conclusion, we identified the brainwide networks underlying fluctuations in delta and alpha rhythms across sleep and wakefulness. Each neural rhythm was captured by distinct aspects of fMRI signals, with alpha predictions driven by two separable networks—one representing the visual system, and the other representing arousal-controlling regions—and delta being predicted by a diffuse array of primarily cortical regions. This work identifies the large-scale activity patterns underlying the neural rhythms that define sleep and wakefulness. Finally, this approach establishes a framework for predictive analyses of brain imaging data that can be employed to identify the underlying network dynamics of diverse neural rhythms.

METHODS

EEG-fMRI acquisition

Both datasets analyzed in this work were collected as part of other studies. Collection for Dataset A was ongoing while our analyses were conducted; the reported number of subjects for this dataset is the number that had been collected at the time our analyses began. Subsequently collected subjects were not counted or considered for the purposes of this study. Dataset B is from previously published studies [13,39].

All subjects reported no neurological, psychiatric, or sleep disorders, and were screened for the following: no MR contraindications, have habitual sleep duration of no less than six hours, daily caffeine consumption of less than 300 mg, not smoke cigarettes, weigh less than 250 pounds, not be pregnant, be comfortable sleeping on their back, and not be taking medications that affect sleep.

Subjects were instructed to lie still inside the scanner with their eyes closed and to allow themselves to fall asleep. In Dataset A, eye closure was asserted with eye tracking. A subset of subjects (17 in Dataset A, 5 in Dataset B) was additionally instructed to press a button on every breath in and out, and to return to performing the task if they awoke during the scan; a lack of button presses indicated subject was asleep [39]. Button press data was not analyzed in this study.

Dataset A

We obtained written informed consent from 22 healthy young adults; 21 (9 male, 12 female; mean age 23.4) were analyzed for the purposes of this study. One subject was excluded due to unusable EEG data caused by excessive subject

motion. The protocol was approved by the Boston University Charles River Campus Institutional Review Board (where 17 subjects were collected) and by the Massachusetts Institute of Technology Institutional Review Board (where 5 subjects were collected). Subjects were scanned on a 3T Siemens Prisma scanner with a 64-channel head-and-neck coil. Anatomical references were acquired using a 1 mm isotropic T1-weighted multi-echo MPRAGE [69]. Functional runs were acquired using TR = 0.378 s, TE = 31 ms, 2.5 mm isotropic voxels, 40 slices, Multiband factor=8, blipped CAIPI shift = FOV/4, flip angle=37°, no in-plane acceleration. Up to three successive runs of 25 minutes each were obtained per subject contingent on subject comfort.

EEG was acquired using MR-compatible 32-channel EEG caps fitted with 4 carbon wire loops (BrainProducts GmbH, Germany) at a sampling rate of 5000 Hz. EEG acquisition was synchronized to the scanner 10 MHz clock to reduce aliasing of high-frequency gradient artifacts. Additional sensors were used to record systemic physiology: respiration was measured simultaneously using an MRI-safe pneumatic respiration transducer belt around the abdomen and pulse was measured with a photoplethysmogram (PPG) transducer (BIOPAC Systems, Inc., Goleta, CA, USA). These systemic physiology signals were not analyzed in this study.

Dataset B

We used a previously published dataset which had collected data from 13 healthy young adults; 11 were analyzed for the purposes of this study (2 male, 9 female, mean age 26.7). One subject was excluded due to a software error that caused fMRI data to be collected at a different temporal resolution, and another was excluded due to poor-quality EEG data caused by excessive subject motion. The protocol was approved by the Massachusetts General Hospital Institutional Review Board. Subjects were scanned on a 3T Siemens Prisma scanner with a 64-channel head-and-neck coil. Anatomical references were acquired using a 1 mm isotropic T1-weighted multi-echo MPRAGE [69]. Functional runs acquired 40 interleaved BOLD-weighted EPI slices with 2.5 mm^3 isotropic voxels. fMRI protocols consisted of a single-shot gradient echo SMS-EPI [70] with MultiBand factor=8, matrix=92×92, blipped CAIPI shift=FOV/4, T*R*=367 ms, TE=30-32 ms (changed midway through the study due to a software upgrade) nominal echo-spacing=0.53 ms, flip angle=32-37°, and no in-plane acceleration. VERSE factor was set between 1 and 1.5 depending on individual subject SAR constraints. Individual runs could last up to 2 hours. If runs ended earlier, subsequent runs would be started up to a maximum total scan duration of 2.5 hours as long as subjects were still comfortable and sleeping.

EEG was acquired using MR-compatible 256-channel geodesic nets and a NA410 amplifier (Electrical Geodesics, Inc., Eugene, OR USA) at a sampling rate of 1000 Hz. EEG acquisition was synchronized to the scanner 10 MHz clock to reduce aliasing of high-frequency gradient artifacts. The scanner cryopump was temporarily shut off during EEG acquisition to reduce vibrational artifact. To acquire reference signals to be used for EEG noise removal [71], subjects wore a reference layer cap composed of an isolating vinyl layer and conductive satin layer on the head, with grommets inserted to allow electrodes to pass through and make contact with the scalp, while other electrodes remained isolated from the scalp and recorded the noise, resulting in a total of 30–36 EEG electrodes per subject. Physiological signals were simultaneously acquired using a Physio16 device (Electrical Geodesics, Inc., Eugene, OR USA). ECG was measured through two disposable electrodes placed on the chest diagonally across the heart, with an MR-compatible lead (InVivo Corp, Philips). Respiration was measured through a piezoelectrical belt (UFI systems, Morro Bay, CA USA) around the chest.

EEG pre-processing

Dataset A

Gradient artifacts were removed using average artifact subtraction [72], with moving average of the previous 20 TRs. Ballistocardiogram artifacts were removed from each EEG channel using signals from the 4 carbon wire loops using the sliding Hanning window regression method from the EEGlab CWL toolbox [73] with the following parameters: window = 25 s and a lag = 0.09 ms. Clean EEG signals were re-referenced to the average of EEG channels.

Dataset B

Gradient artifacts were removed using average artifact subtraction [72], with moving average of the previous 20 TRs. Electrodes were then re-referenced to the common average, computing this separately for electrodes contacting the

head, and those placed on the reference layer. Channels on the cheeks and borders of the reference cap were excluded from the common average. Ballistocardiogram artifacts were removed using regression of reference signals from the isolated EEG electrodes [74]. Since there was a larger number of noise electrodes than signal electrodes, the regression was performed after subsampling the noise electrodes, using only every fourth isolated electrode. Because the position of and physiological noise influences on the electrodes can vary over the long recording times used here, we implemented a dynamic time-varying regression of the reference signals. Beta coefficients for the best-fit regression within 30 s sliding time windows were fit using least-squares; these beta values were then linearly interpolated over the nonoverlapping windows. The resulting interpolated beta value at every time point was then used for a local subtraction of the reference signals from the modeled EEG recording. This regression was performed individually for each EEG channel.

fMRI pre-processing

The cortical surface was reconstructed from the MEMPRAGE volume using Freesurfer version 6 [75]. All functional runs were slice-time corrected using FSL version 6 (slicetimer; https://fsl.fmrib.ox.ac.uk/fsl/fslwiki) [76] and motion corrected to the middle frame using AFNI (3dvolreg; https://afni.nimh.nih.gov/). Each motion-corrected run was then registered to the anatomical data using boundary-based registration (bbregister) [77].

For Dataset A, physiological noise was removed using HRAN, open-source software implementing a statistical model of harmonic regression with autoregressive noise [78].

For Dataset B, Physiological noise removal was performed using dynamic regression based on the concept of RETROICOR [79] and adapted for fast fMRI as follows. The respiratory trace was bandpass filtered between 0.16–0.4 Hz using a finite impulse response filter and the instantaneous phase was computed as the angle of the Hilbert transform. The cardiac peaks were detected automatically using the FASST toolbox (montefiore.ulg.ac.be/~phillips/FASST) and the phase was modeled as varying linearly between each identified peak. Sine and cosine basis functions using the phase of the signal and its second harmonic were generated as regressors for physiological noise. This regression was performed over 1000 s windows sliding every 400 s to enable high-quality physiological noise removal as the heart rate and respiratory rate varied throughout the scan. No spatial smoothing was applied.

Data preparation

Data preparation, model training, and statistical analyses were conducted in MATLAB unless noted otherwise. EEG power was extracted from occipital electrodes via the multitaper spectrogram method [80,81]. Spectrograms were obtained using a 5 second window, sliding every 0.1 s. Spectrograms were visually inspected for the presence of clear eyes-closed alpha rhythms. Most healthy adults display a clear increase in alpha power when awake with eyes closed (see Fig. 1a), but a small percentage of individuals do not show this increase for reasons not yet understood [36,37,82,83]. As this difference may reflect distinct underlying networks, we excluded these subjects from the alpha analyses (3 subject in Dataset A, 5 subjects in Dataset B). These subjects were included in the delta analyses.

The power time series was then extracted for the frequencies of interest: 1.1-4 Hz for delta (chosen to minimize contamination from low-frequency artifacts in EEG-fMRI) and 8.5-12 Hz for alpha (chosen to minimize contamination from theta and spindle frequency ranges while still covering subject variations in alpha range). Power in the 40-52 Hz range was also extracted for the purpose of artifact detection (see below).

For dataset A, power was extracted from three occipital electrodes (O1, O2, and Oz) and then averaged. For dataset B, which used a reference layer for artifact removal [71] and therefore did not have identical channel layouts across subjects, power was extracted from the single occipital electrode closest to Oz which had good data quality based on visual inspection of the spectrogram. This inspection procedure was done for a prior study [13] and we used the same previously selected electrodes. Power time series were then interpolated to match fMRI TR timings.

fMRI segmentation was performed in FreeSurfer based on the Desikan-Killiany atlas [84]. The following regions were obtained: 31 bilateral cortical regions (the entire cortex minus the temporal pole, frontal pole, and entorhinal regions, which were not covered by the acquisition volume in all subjects), 7 bilateral subcortical regions (thalamus, pallidum,

amygdala, caudate, putamen, accumbens, hippocampus), and 8 non-grey matter regions (left- and right- total white matter as an index of global blood flow, left- and right- lateral and inferior lateral ventricles, third ventricle, and fourth ventricle); total of 84 regions. Anatomical labels were interpolated to functional space, and functional voxels with 70% or above of overlap with the label were averaged to generate the mean region time series. The first 20 TRs of each run were discarded, and the fMRI timeseries of each region (separately for each subject) was z-scored.

To mark data points unsuitable for analyses, the following criteria were used: 1) TRs with motion above 0.3 mm; 2) z-scored EEG power in the 40-52Hz range with a value above 1 (adjusted to be lower for runs with particularly noisy data); 3) detected local outliers (distance larger than 6 local scaled median absolute deviation (MAD)) within a sliding window of 1000 TRs (for dataset B, a threshold of 3 MAD was used instead due to noisier data; thresholds were also stricter for subjects with excessively noisy data). The outlier detection method was applied to delta power and power in the 40-52Hz range (but not to alpha power, as its naturally fast fluctuations would be erroneously detected as outliers; however, three subjects in Dataset A with excessively noisy data underwent alpha band outlier removal based on a smaller window). Datapoints that fulfilled any of these criteria were then removed, and the remaining data was saved as separate segments to avoid creation of spurious continuities (segments below 60 TRs were discarded, as 60 TRs of continuous fMRI are required to predict each EEG point). EEG data for each subject was then rescaled to 0-1.

Model training and testing

Following the above, fMRI data was split into overlapping 60-TR sections, with the goal of predicting the single EEG power point at the center of the fMRI window (equivalent to the position of the 30th fMRI point). We predicted the EEG point at the center of the window to avoid assumptions about temporal relationship between the two modalities (such as regions which might have early or late responses), as our goal in this work was to uncover spatial patterns. Early model explorations revealed similar prediction performance for segment sizes in the 40 to 80 TR range (roughly 14s to 29s); 60 TRs was ultimately chosen for being long enough to include the width of the canonical HRF in the post-EEG half of the segment, while being short enough to avoid excessively inflating the number of weights that the model had to learn.

In order to predict every EEG point that had been interpolated to match the fMRI TR timings, there was a 59-TR overlap between each adjacent predictor window (windows were slid by 1 TR). Given that no fMRI points would be available for EEG points at the edges of the data segments following artifact removal, the first 29 and the last 30 EEG points of each segment were discarded.

The model was trained on a maximum of 5040 predictors (60 TR segments and 84 parcellated regions; region list under 'Data preparation'). Given that different analyses used different numbers of parcellated regions, the actual number of predictors varied by analysis (i.e. for the results presented in Fig. 3a, the model was trained on each bilateral region at a time; thus, only 120 predictors).

At each training fold, all data for one subject was held-out as test data. The model was then trained on the remaining subjects using the Matlab *fitrlinear* function. The function fits a weight value for each input feature (each TR of each input fMRI region), along with a bias value, and a lambda regularization term. The regularization term multiples the sum of squares of all model weights to penalize large weights and prevent the model from overfitting (L2 regularization). During training, the model adjusts its weights and bias using stochastic gradient descent.

In each fold, the lambda regularization term was adjusted through Bayesian hyperparameter optimization by partitioning the training data (i.e. the held-out subject was not used) in 5 folds over 30 iterations, with repartitioning at every iteration. The best lambda value was then selected, and the model was trained on all training data at once. While the model frequently chose very similar lambda values in each held-out subject fold, this additional computational effort resulted in a strict training regimen where no data from the held-out subject was ever used to tune the model. In earlier explorations of the model, the Bayesian hyperparameter optimization was also allowed to select between L1 and L2 regularization, but given that L2 consistently outperformed L1, we fixed L2 as the regularization option to reduce computational processing time.

Once the optimal lambda value was determined and the weights and biases were fit on the training data, the model was fixed and used to generate predictions on the held-out subject. We highlight that normalization (see 'Data preparation' above) was done separately within each subject so that at no point did the held-out subject data affect training data. A Pearson's linear correlation coefficient was then calculated between the predictions made for the held-out subject and the ground truth.

In order to achieve more precise estimates of predictive performance, we repeated each train-and-test fold 10 times with distinct random seeds, yielding 10 trained models and thus 10 sets of predictions (plus 10 Pearson's correlation values) for each held-out subject. In all analyses, each correlation value reported for each held-out subject represents the mean correlation performance of 10 models.

Statistical analyses

Major Anatomical groups analysis

The model underwent the training and testing routine described above using different groups of fMRI regions as input: only the subcortical regions, only the cortical regions, only the non-gray matter regions, and all of them together. Finally, the model also went through the same training routine using all the regions as input, but with fMRI data shuffled in relation to the EEG (which was kept unshuffled in order to preserve its original temporal structure). Once the correlation performance was obtained for all subjects and conditions, an ANOVA was conducted, with subject ID as a random variable, and condition as a fixed variable. Pairwise comparisons between each condition were then performed using Tukey's HSD (honestly significant difference, two-sided), conducted within each dataset (A/B/combined) and analyses type (alpha/delta). Tukey's HSD corrects for multiple comparisons and operates upon the ANOVA results [85].

Bootstrapping for single-region analysis

To assess whether prediction performance differences between individual regions were not the result of the ordering of random noise, we shuffled the region labels within each subject, then obtained region means, ordered the means, and fit a line to the ordered means. This was repeated 1000 times; the resulting lines were then averaged, and a 95% confidence interval (based on number of subjects) was calculated. Verification that the line fit to the real (non-shuffled) data was outside these bounds was taken to mean that region differences (i.e. spatial pattern of predictive performance across fMRI regions) were significant and not due to ordering of random noise. Following this analysis, paired two-tailed t-tests were conducted to assess which regions were superior to shuffled control (or to non-gray matter regions in the case of the delta analyses).

Clustering of pairwise performance

The model was trained and tested on each pair of gray matter regions. For delta, the non-gray matter regions were also added to each gray matter pair. For each pair, the performance values from the single-region analysis (correlation between predictions and ground truth) were noted, and the highest performance was subtracted from the pair's performance. This value represented the performance gain (or loss) of the pair in relation to its best individual component. These values were then used for clustering.

To assess the optimal clustering solution (including the null hypothesis of no suitable clustering), we used the gap statistic [44]. The gap value is defined as the difference between the within-cluster dispersion obtained by Monte Carlo sampling from a reference distribution (generated from a uniform distribution over a box aligned with the principal components of the data) and the actual within-cluster dispersion obtained by clustering the real data. The ideal number of clusters is defined as the smallest number of clusters that satisfies $Gap(k) \geq Gap(k+1) - SE(k+1)$, where $Gap(k)$ and $Gap(k+1)$ are the gap values for the clustering solutions with k and k+1 clusters respectively, and $SE(k+1)$ is the standard error of the clustering solution with k+1 clusters.

Spatial scales analysis

After training and testing of the model using all 84 parcellated regions, the model's beta weights were extracted, and weights (across all TRs) for each region were squared and summed. The regions that displayed weights whose sum of squares were below the bottom 10% (as distribution percentiles) were then removed from the model, which was then

once again trained and tested. This was repeated up until 5 regions were left in the model. Performance in each iteration was then compared to performance using all brain regions with two-tailed t-tests corrected using the Benjamini-Hochberg procedure [86].

ACKNOWLEDGEMENTS

This study was funded by National Institutes of Health grants R01-AG070135, U19-NS128613, and U19-NS123717; the Sloan Fellowship, the McKnight Scholar Award, the Pew Biomedical Scholar Award, the Simons Collaboration on Plasticity in the Aging Brain (811231), and the One Mind Rising Star Award. Resources were provided by NSF Major Research Instrumentation grant BCS-1625552. Collaboration was facilitated by the Janelia Visiting Scientist Program.

Author contributions

L.P.L.J. and L.D.L. conceptualized the study. S.M.B, S.D.W., and L.D.L. designed the experiments. S.M.B and S.D.W. collected and pre-processed the data, and wrote the corresponding sections of the methods. L.P.L.J. conducted data preparation, performed the analyses, created the figures, and wrote the manuscript. C.S. advised on modeling and analytical methods. C.S. and L.D.L. revised the manuscript and all authors provided edits.

Competing interests

The authors declare no competing interests.

REFERENCES

1. Adamantidis, A. R., Gutierrez Herrera, C. & Gent, T. C. Oscillating circuitries in the sleeping brain. *Nat Rev Neurosci* **20**, 746–762 (2019).

2. Geva-Sagiv, M. *et al.* Augmenting hippocampal–prefrontal neuronal synchrony during sleep enhances memory consolidation in humans. *Nat Neurosci* **26**, 1100–1110 (2023).

3. Helfrich, R. F., Mander, B. A., Jagust, W. J., Knight, R. T. & Walker, M. P. Old Brains Come Uncoupled in Sleep: Slow Wave-Spindle Synchrony, Brain Atrophy, and Forgetting. *Neuron* **97**, 221-230.e4 (2018).

4. Holz, J. *et al.* EEG sigma and slow-wave activity during NREM sleep correlate with overnight declarative and procedural memory consolidation. *J Sleep Res* **21**, 612–619 (2012).

5. Huber, R., Ghilardi, M. F., Massimini, M. & Tononi, G. Local sleep and learning. **430**, 4–7 (2004).

6. Latchoumane, C. F. V., Ngo, H. V. V., Born, J. & Shin, H. S. Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and Hippocampal Rhythms. *Neuron* **95**, 424-435.e6 (2017).

7. Maingret, N., Girardeau, G., Todorova, R., Goutierre, M. & Zugaro, M. Hippocampo-cortical coupling mediates memory consolidation during sleep. *Nat Neurosci* **19**, 959–964 (2016).

8. Marshall, L., Helgadóttir, H., Mölle, M. & Born, J. Boosting slow oscillations during sleep potentiates memory. *Nature* **444**, 610–613 (2006).

9. Ngo, H. V. V., Martinetz, T., Born, J. & Mölle, M. Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. *Neuron* **78**, 545–553 (2013).

10. Schreiner, T., Petzka, M., Staudigl, T. & Staresina, B. P. Endogenous memory reactivation during sleep in humans is clocked by slow oscillation-spindle complexes. *Nat Commun* **12**, (2021).

11. Benwell, C. S. Y. *et al.* Frequency and power of human alpha oscillations drift systematically with time-on-task. *Neuroimage* **192**, 101–114 (2019).

12. Schneider, D., Herbst, S. K., Klatt, L. I. & Wöstmann, M. Target enhancement or distractor suppression? Functionally distinct alpha oscillations form the basis of attention. *European Journal of Neuroscience* **55**, 3256–3265 (2022).

13. Fultz, N. E. *et al.* Coupled electrophysiological, hemodynamic, and cerebrospinal fluid oscillations in human sleep. *Science (1979)* **366**, 628–631 (2019).

14. Pfurtscheller, G., Stancák, A. & Neuper, Ch. Event-related synchronization (ERS) in the alpha band — an electrophysiological correlate of cortical idling: A review. *International Journal of Psychophysiology* **24**, 39–46 (1996).

15. Clayton, M. S., Yeung, N. & Cohen Kadosh, R. The many characters of visual alpha oscillations. *European Journal of Neuroscience* **48**, 2498–2508 (2018).

16. Xie, L. *et al.* Sleep Drives Metabolite Clearance from the Adult Brain. *Science (1979)* **342**, 373–377 (2013).

17. Martínez-Montes, E., Valdés-Sosa, P. A., Miwakeichi, F., Goldman, R. I. & Cohen, M. S. Concurrent EEG/fMRI analysis by multiway Partial Least Squares. *Neuroimage* **22**, 1023–1034 (2004).

18. de Munck, J. C. *et al.* The hemodynamic response of the alpha rhythm: An EEG/fMRI study. *Neuroimage* **35**, 1142–1151 (2007).

19. Feige, B. *et al.* Cortical and subcortical correlates of electroencephalographic alpha rhythm modulation. *J Neurophysiol* **93**, 2864–2872 (2005).

20. Goldman, R., Stern, J., Engel, J. & Cohen, M. S. Simultaneous EEG and fMRI of the alpha rhythm. *Neuroreport* 2487–2492 (2002) doi:10.1097/01.wnr.0000047685.08940.d0.

21. Tyvaert, L., LeVan, P., Grova, C., Dubeau, F. & Gotman, J. Effects of fluctuating physiological rhythms during prolonged EEG-fMRI studies. *Clinical Neurophysiology* **119**, 2762–2774 (2008).

22. Moosmann, M. *et al.* Correlates of alpha rhythm in functional magnetic resonance imaging and near infrared spectroscopy. *Neuroimage* **20**, 145–158 (2003).

23. Laufs, H. *et al.* EEG-correlated fMRI of human alpha activity. *Neuroimage* **19**, 1463–1476 (2003).

24. Liu, Z. *et al.* Finding thalamic BOLD correlates to posterior alpha EEG. *Neuroimage* **63**, 1060–1069 (2012).

25. Dang-Vu, T. T. *et al.* Cerebral correlates of delta waves during non-REM sleep revisited. *Neuroimage* **28**, 14–21 (2005).

26. Handwerker, D. A., Ollinger, J. M. & D'Esposito, M. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage* **21**, 1639–1651 (2004).

27. Lewis, L. D., Setsompop, K., Rosen, B. R. & Polimeni, J. R. Fast fMRI can detect oscillatory neural activity in humans. *Proceedings of the National Academy of Sciences* **113**, E6679–E6685 (2016).

28. Polimeni, J. R. & Lewis, L. D. Imaging faster neural dynamics with fast fMRI: A need for updated models of the hemodynamic response. *Progress in Neurobiology* vol. 207 Preprint at https://doi.org/10.1016/j.pneurobio.2021.102174 (2021).

29. Lindquist, M. A., Meng Loh, J., Atlas, L. Y. & Wager, T. D. Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling. *Neuroimage* **45**, S187–S198 (2009).

30. Shen, X. *et al.* Using connectome-based predictive modeling to predict individual behavior from brain connectivity. *Nat Protoc* **12**, 506–518 (2017).

31. Huth, A. G., De Heer, W. A., Griffiths, T. L., Theunissen, F. E. & Gallant, J. L. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* **532**, 453–458 (2016).

32. Finn, E. S. & Rosenberg, M. D. Beyond fingerprinting: Choosing predictive connectomes over reliable connectomes. *Neuroimage* **239**, 118254 (2021).

33. Lake, E. M. R. *et al.* Simultaneous cortex-wide fluorescence Ca2+ imaging and whole-brain fMRI. *Nat Methods* **17**, 1262–1271 (2020).

34. Chang, C. *et al.* Tracking brain arousal fluctuations with fMRI. *Proc Natl Acad Sci U S A* **113**, 4518–4523 (2016).

35. Falahpour, M., Chang, C., Wong, C. W. & Liu, T. T. Template-based prediction of vigilance fluctuations in resting-state fMRI. *Neuroimage* **174**, 317–327 (2018).

36. Davis, H., Davis, P. A., Loomis, A. L., Harvey, E. N. & Hobart, G. HUMAN BRAIN POTENTIALS DURING THE ONSET OF SLEEP. *J Neurophysiol* **1**, 24–38 (1938).

37. DAVIS, H. ACTION POTENTIALS OF THE BRAIN. *Arch Neurol Psychiatry* **36**, 1214 (1936).

38. Noesselt, T. *et al.* Audiovisual temporal correspondence modulates human multisensory superior temporal sulcus plus primary sensory cortices. *Journal of Neuroscience* **27**, 11431–11441 (2007).

39. Setzer, B. *et al.* A temporal sequence of thalamic activity unfolds at transitions in behavioral arousal state. *Nat Commun* **13**, (2022).

40. Lindquist, M. A., Geuter, S., Wager, T. D. & Caffo, B. S. Modular preprocessing pipelines can reintroduce artifacts into fMRI data. *Hum Brain Mapp* **40**, 2358–2376 (2019).

41. Chen, J. E., Jahanian, H. & Glover, G. H. Nuisance Regression of High-Frequency Functional Magnetic Resonance Imaging Data: Denoising Can Be Noisy. *Brain Connect* **7**, 13–24 (2017).

42. Bright, M. G., Tench, C. R. & Murphy, K. Potential pitfalls when denoising resting state fMRI data using nuisance regression. *Neuroimage* **154**, 159–168 (2017).

43. Murphy, M. *et al.* Source modeling sleep slow waves. *Proc Natl Acad Sci U S A* **106**, 1608–1613 (2009).

44. Tibshirani, R., Walther, G. & Hastie, T. Estimating the number of clusters in a data set via the gap statistic. *J R Stat Soc Series B Stat Methodol* **63**, 411–423 (2001).

45. Lazarus, M., Chen, J. F., Urade, Y. & Huang, Z. L. Role of the basal ganglia in the control of sleep and wakefulness. *Curr Opin Neurobiol* **23**, 780–785 (2013).

46. Haber, S. N. & Calzavara, R. The cortico-basal ganglia integrative network: The role of the thalamus. *Brain Res Bull* **78**, 69–74 (2009).

47. Vogt, B. A., Pandya, D. N. & Rosene, D. L. Cingulate cortex of the rhesus monkey: I. Cytoarchitecture and thalamic afferents. *Journal of Comparative Neurology* **262**, 256–270 (1987).

48. Jang, S. H. & Yeo, S. S. Thalamocortical tract between anterior thalamic nuclei and cingulate gyrus in the human brain: Diffusion tensor tractography study. *Brain Imaging Behav* **7**, 236–241 (2013).

49. Silva, L. R., Amitai, Y. & Connors, B. W. Intrinsic Oscillations of Neocortex Generated by Layer 5 Pyramidal Neurons. *Science (1979)* **251**, 432–435 (1991).

50. Zumer, J. M., Scheeringa, R., Schoffelen, J. M., Norris, D. G. & Jensen, O. Occipital Alpha Activity during Stimulus Processing Gates the Information Flow to Object-Selective Cortex. *PLoS Biol* **12**, (2014).

51. Betta, M. *et al.* Cortical and subcortical hemodynamic changes during sleep slow waves in human light sleep. *Neuroimage* **236**, 118117 (2021).

52. Ilhan-Bayrakcı, M., Cabral-Calderin, Y., Bergmann, T. O., Tüscher, O. & Stroh, A. Individual slow wave events give rise to macroscopic fMRI signatures and drive the strength of the BOLD signal in human resting-state EEG-fMRI recordings. *Cereb Cortex* **32**, 4782–4796 (2022).

53. Dang-Vu, T. T. *et al.* Spontaneous neural activity during human slow wave sleep. *Proc Natl Acad Sci U S A* **105**, 15160–15165 (2008).

54. Raut, R. V *et al.* Global waves synchronize the brain's functional systems with fluctuating arousal. *Sci Adv* **7**, 1–16 (2021).

55. Mitra, A. *et al.* Human cortical-hippocampal dialogue in wake and slow-wave sleep. *Proc Natl Acad Sci U S A* **113**, E6868–E6876 (2016).

56. Massimini, M., Huber, R., Ferrarelli, F., Hill, S. & Tononi, G. The sleep slow oscillation as a traveling wave. *Journal of Neuroscience* **24**, 6862–6870 (2004).

57. Malerba, P., Whitehurst, L. N., Simons, S. B. & Mednick, S. C. Spatio-temporal structure of sleep slow oscillations on the electrode manifold and its relation to spindles. *Sleep* **42**, 1–14 (2019).

58. Bolt, T. *et al.* A parsimonious description of global functional brain organization in three spatiotemporal patterns. *Nat Neurosci* **25**, 1093–1103 (2022).

59. Shine, J. M. *et al.* Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nat Neurosci* **22**, 289–296 (2019).

60. Saletin, J. M., van der Helm, E. & Walker, M. P. Structural brain correlates of human sleep oscillations. *Neuroimage* **83**, 658–668 (2013).

61. Li, W., Sun, L., Yue, L. & Xiao, S. The left temporal transverse cortex is affected by poor sleep quality, which in turn contributes to depressive symptoms in older adults. *Heliyon* **9**, (2023).

62. Jacob, L. P. L. & Huber, D. E. Neural Habituation Enhances Novelty Detection: an EEG Study of Rapidly Presented Words. *Comput Brain Behav* **3**, 208–227 (2020).

63. Sarma, A., Masse, N. Y., Wang, X.-J. & Freedman, D. J. Task-specific versus generalized mnemonic representations in parietal and prefrontal cortices. *Nat Neurosci* **19**, 143–149 (2016).

64. Spellman, T. *et al.* Hippocampal–prefrontal input supports spatial encoding in working memory. *Nature* **522**, 309–314 (2015).

65. Chang, L. J., Gianaros, P. J., Manuck, S. B., Krishnan, A. & Wager, T. D. A sensitive and specific neural signature for picture-induced negative affect. *PLoS Biol* **13**, 1–28 (2015).

66. Meir-Hasson, Y., Kinreich, S., Podlipsky, I., Hendler, T. & Intrator, N. An EEG Finger-Print of fMRI deep regional activation. *Neuroimage* **102**, 128–141 (2014).

67. Simões, M. *et al.* How much of the BOLD-fMRI signal can be approximated from simultaneous EEG data: Relevance for the transfer and dissemination of neurofeedback interventions. *J Neural Eng* **17**, (2020).

68. Kucyi, A. *et al.* Intracranial electrophysiology reveals reproducible intrinsic functional connectivity within human brain networks. *Journal of Neuroscience* **38**, 4230–4242 (2018).

69. van der Kouwe, A. J. W., Benner, T., Salat, D. H. & Fischl, B. Brain morphometry with multiecho MPRAGE. *Neuroimage* **40**, 559–569 (2008).

70. Setsompop, K. *et al.* Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn Reson Med* **67**, 1210–1224 (2012).

71. Luo, Q., Huang, X. & Glover, G. H. Ballistocardiogram artifact removal with a reference layer and standard EEG cap. *J Neurosci Methods* **233**, 137–149 (2014).

72. Allen, P. J., Josephs, O. & Turner, R. A Method for Removing Imaging Artifact from Continuous EEG Recorded during Functional MRI. *Neuroimage* **12**, 230–239 (2000).

73. van der Meer, J. N. *et al.* Carbon-wire loop based artifact correction outperforms post-processing EEG/fMRI corrections—A validation of a real-time simultaneous EEG/fMRI correction method. *Neuroimage* **125**, 880–894 (2016).

74. Chowdhury, M. E. H., Mullinger, K. J., Glover, P. & Bowtell, R. Reference layer artefact subtraction (RLAS): A novel method of minimizing EEG artefacts during simultaneous fMRI. *Neuroimage* **84**, 307–319 (2014).

75. Fischl, B. FreeSurfer. *Neuroimage* **62**, 774–781 (2012).

76. Woolrich, M. W. *et al.* Bayesian analysis of neuroimaging data in FSL. *Neuroimage* **45**, S173–S186 (2009).

77. Greve, D. N. & Fischl, B. Accurate and robust brain image alignment using boundary-based registration. *Neuroimage* **48**, 63–72 (2009).

78. Agrawal, U., Brown, E. N. & Lewis, L. D. Model-based physiological noise removal in fast fMRI. *Neuroimage* **205**, (2020).

79. Glover, G. H., Li, T. Q. & Ress, D. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn Reson Med* **44**, 162–167 (2000).

80. Thomson, D. J. Spectrum estimation and harmonic analysis. *Proceedings of the IEEE* **70**, 1055–1096 (1982).

81. Prerau, M. J., Brown, R. E., Bianchi, M. T., Ellenbogen, J. M. & Purdon, P. L. Sleep Neurophysiological Dynamics Through the Lens of Multitaper Spectral Analysis. (2017) doi:10.1152/physiol.00062.2015.

82. Barzegaran, E., Vildavski, V. Y. & Knyazeva, M. G. Fine Structure of Posterior Alpha Rhythm in Human EEG: Frequency Components, Their Cortical Sources, and Temporal Behavior. *Sci Rep* **7**, 1–12 (2017).

83. Niedermeyer, E. The Normal EEG of the Waking Adult. in *Electroencephalography: Basic Principles, Clinical Applications and Related Fields* 149–173 (Williams and Wilkins, Philadelphia, 1999).

84. Desikan, R. S. *et al.* An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* **31**, 968–980 (2006).

85. Tukey, J. W. Comparing individual means in the analysis of variance. *Biometrics* **5**, 99–114 (1949).

86. Benjamini, Y. & Hochberg, Y. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *J R Stat Soc Series B Stat Methodol* **57**, 289–300 (1995).