Uncovering dynamic human brain phase coherence networks

Anders S. Olsen[1*], Anders Brammer[1], Patrick M. Fisher[2,3], Morten Mørup[1*]

[1]Department of Applied Mathematics and Computer Science, Technical University of Denmark, Kgs. Lyngby, Denmark

[2]Neurobiology Research Unit, Copenhagen University Hospital Rigshospitalet, Copenhagen, Denmark

[3]Department of Drug Design and Pharmacology, Faculty of Health and Medical Sciences, University of Copenhagen, Copenhagen, Denmark

[*]Corresponding authors: ansol@dtu.dk, mmor@dtu.dk

Keywords: Complex angular central Gaussian, Dynamic functional connectivity, Probabilistic mixture modeling, Phase coherence, Multivariate synchronization.

Significance statement

Understanding how the human brain coordinates activity across distant regions is central to explaining cognition and behavior. Most existing approaches study these interactions by tracking changes in signal strength, which can be strongly affected by non-neural artifacts. Here, we focus instead on the *phase relationships* between brain signals: How brain regions synchronize their oscillations forming dynamic phase coherence networks. We introduce a flexible and principled mixture modeling framework to capture these patterns directly and reveal recurring states of brain-wide synchronized activity consistent across individuals. This approach uncovers meaningful differences between diverse cognitive tasks, without requiring task labels during training. By emphasizing signal phase rather than amplitude, our method offers a robust and interpretable way to study dynamic brain organization.

ABSTRACT

Complex cognitive functions rely on coordinated communication between distributed brain regions, yet capturing these interactions as they evolve over time remains challenging. Traditional analyses of functional brain connectivity largely rely on correlations in signal amplitude, which are sensitive to noise and artifacts such as head motion. Here, we introduce a mixture modeling approach that focuses on the phase of brain signals, allowing dynamic patterns of large-scale synchronization in brain phase coherence networks to be studied directly and in their entirety. We lay the mathematical and conceptual groundwork for phase modeling and introduce the complex angular central Gaussian distribution, providing a principled way to analyze phase-based interactions across the brain. Applied to fMRI data, the model identifies recurring states of brain-wide synchronized activity that reliably distinguish cognitive tasks and generalize across previously unseen individuals, without requiring any task labels during training. These results show that modeling signal phase offers a cleaner and more informative view of brain dynamics than amplitude-based approaches alone, opening new avenues for studying large-scale neural coordination.

INTRODUCTION

Synchronization is a fundamental concept in science, observed across diverse fields from coherent light in lasers, anti-coherence in acoustic noise suppression, and the circadian rhythm aligning with the dark-light cycle of the environment. In neuroscience, synchronization is crucial across multiple scales: Neuron-to-neuron communication often occurs through synchronized bursts of firing [1,2] and synchronized mesoscale activity within cortical neuronal columns generates electrical signals that can be detected on the scalp [3]. At the macroscale, synchronization underpins functional connectivity across brain areas, providing a framework for understanding cognitive functions and behaviors [4–6]. As such, accurate models of phase coherence, the quantification of synchronization, in multi-channel brain data are essential to capture the complexity of brain-wide connectivity patterns, key to understanding the orchestrated dynamics within and across neural networks.

Traditionally, connectivity between regional neuroimaging signals such as those from blood oxygen level dependent functional magnetic resonance imaging (BOLD fMRI) has been measured using the Pearson correlation coefficient [7], which captures the linear relationship between signals [8,9]. Estimating time-varying correlation coefficients requires temporal windowing of input signals, potentially decreasing

stationarity of low-frequency components and leading to challenges in determining, e.g., the optimal window length, function, and stride [10–13], and has been suggested to be an unstable approach to identifying dynamic functional connectivity [14–17], see [18] for a recent review. Directly modeling the time-series using Gaussian mixture models (GMM) or hidden Markov models offer window-free alternatives [19], however, intensity-based time-series methods are vulnerable to amplitude-related noise such as head motion [20]. Instead, assessing phase coherence (a.k.a. phase coupling) between brain signals, focusing exclusively on co-oscillation and discarding (co-)amplitude information, offers a frame-wise connectivity proxy free from area-specific amplitude-related noise uninformative of connectivity [21].

Phase estimation involves converting signals to the complex domain, typically via the Hilbert transform, which translates the temporal evolution of a signal $x(t)$ into a revolution around the unit circle $z(t) = a(t)\,\mathrm{e}^{i\theta(t)}$ (Fig. 1A). The phase $\theta(t)$ is the angle to the real axis, and co-oscillation between regions is reflected by phase differences assessed at each data sample, i.e., $\theta_j(t) - \theta_{j'}(t)$ for regions j and j'. Since the phase and phase difference is a circular quantity with period 2π and therefore difficult to model arithmetically, most previous fMRI-studies involving phase coherence have focused exclusively on cosine phase coherence, i.e., $\cos\left(\theta_j(t) - \theta_{j'}(t)\right)$ [21–25], which discards the imaginary dimension (the sine), which holds potentially valuable information about phase alignment. In contrast, the sine is often favored in electroencephalography data due to volume conduction [26]. Furthermore, some studies have focused exclusively on the leading eigenvector of the cosine phase coherence matrix, known as Leading Eigenvector Dynamics Analysis (LEiDA) [22], further explicitly discarding a portion of the data variance within the cosine projection of the phase coherence.

In this paper, we show that interregional phase coherence is very easily modeled when considering the full rectangular form $e^{i(\theta_j(t)-\theta_{j'}(t))}$. The complex dynamic phase coherence map between p regions $\mathbf{A}_t \in \mathbb{C}^{p\times p}$ for each volume t is rank-1, i.e., $\mathbf{A}_t = e^{i\boldsymbol{\theta}_t}\left(e^{i\boldsymbol{\theta}_t}\right)^{\mathrm{H}}$, where $(\cdot)^{\mathrm{H}}$ is the complex conjugate transpose, rendering modeling a pool of instantaneous complex-valued phase vectors $\left\{\mathbf{u}_t = e^{i\boldsymbol{\theta}_t}\right\}_{t=1}^{N}$ straightforward without the need for dimensionality reduction. Importantly, such phase vectors are naturally unit norm and sign-invariant and therefore reside on the complex projective hyperplane $\mathbb{CP}^{p-1}$. The Gaussian analogue on this manifold is the complex angular central Gaussian (CACG) distribution (Methods) [27,28], which is parameterized by a full covariance matrix. However, macroscale brain function has routinely been posited to be low rank [29,30] and in high spatial resolution the number of estimated pairwise covariances increases quadratically. Therefore, a rank-r reparameterization, where $r \leq p$, can

be applied to explicitly control the number of estimated parameters, deriving a more parsimonious and interpretable representation of brain connectivity.

Here we propose a new dynamic modeling framework for the detailed characterization of phase coherence networks using probabilistic mixture modeling using CACG distributions as mixture components (Fig. 1B). We contrast our interregional phase coherence modeling with traditional time-series modeling via the Gaussian distribution, phase-amplitude modeling via the complex Gaussian distribution, and existing K-means approaches for phase coherence, including LEiDA (see Supporting Table S1 for an overview of all models). All mixture components are parameterized by a covariance matrix, the rank of which we show is important for characterizing the complexity of dynamic states in the data. We apply our mixture models, which are trained without the use of labels, to task fMRI data from the Human Connectome project [31], showcasing how the CACG mixture model learns meaningful phase coherence components reflective of cognitive states. When applied to resting-state fMRI, the model recovers relatively few, complex network patterns. In short, phase modeling with CACG mixtures offers a principled, robust, and interpretable unsupervised toolkit for mapping dynamic brain phase coherence networks. All models have been implemented in our openly available Python toolbox "Phase Coherence Mixture Modeling" (PCMM, github.com/anders-s-olsen/PCMM).

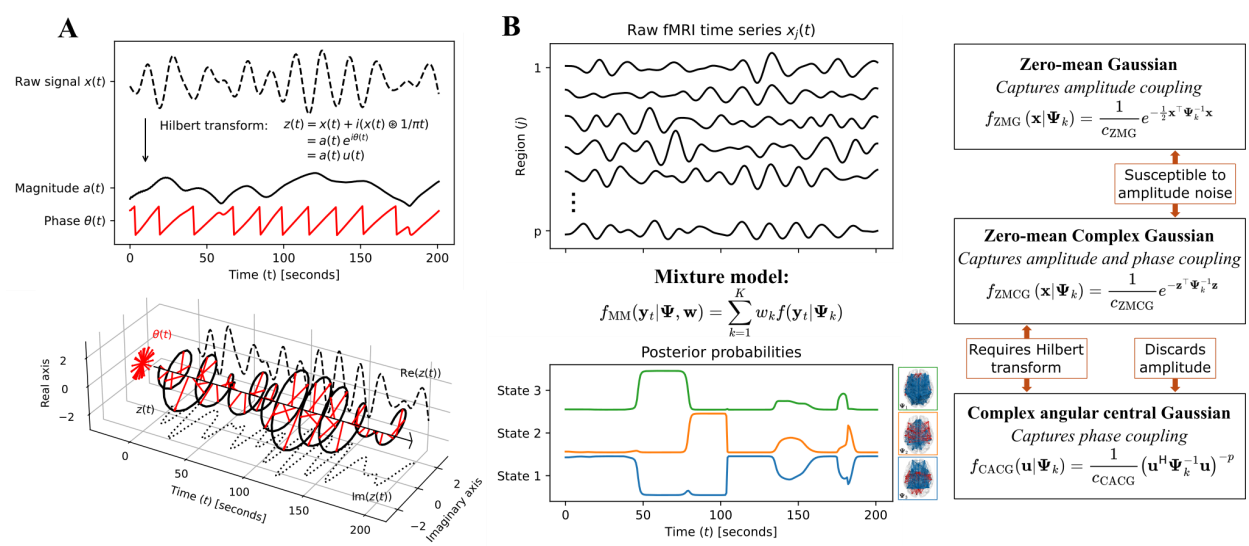

Figure 1: Mixture modeling of fMRI signals using either their raw values, phase or phase-amplitude representations. ***A)****: The Hilbert transform can be used to create the analytic signal* $z(t)$*, which decomposes a time-series* $x(t)$ *into phase* $\theta(t)$ *and amplitude* $a(t)$*. The rectangular representation of phase, discarding amplitude, is* $u(t) = e^{i\theta(t)}$*.* ***B)*** *Mixture modeling of fMRI data can be approached*

using either their Euclidean representation (zero-mean Gaussian), or through the Hilbert transform providing a phase-amplitude representation (complex zero-mean Gaussian). Alternatively, the amplitude can be discarded completely and the phase modeled using the complex angular central Gaussian (CACG) distribution. $^{C}(\cdot)$ *is a normalization constant (see Methods for full densities).*

RESULTS

We first establish the advantage of modeling phase coherence in its entirety using the complex domain that accounts for both the cosine and sine in order to adequately recover synthetic data with systematic phase coherence fingerprints. We subsequently introduce the CACG distribution as the Gaussian analogue for multivariate phase data and fit a mixture model with low-rank CACG distributions as components to task-based and resting-state fMRI data from 255 unrelated subjects from the Human Connectome Project downsampled to 116 cortical and subcortical regions [32,33], showcasing the impact of component rank on model estimates and task classification accuracy, and the benefits of phase coherence modeling compared to amplitude- and phase-amplitude modeling.

The complex domain is required for modeling phase coherence in its entirety

To estimate phase coherence, signals are first converted to the complex domain via the Hilbert transform (Fig. 1A), wherein the amplitude is discarded and only the complex-valued instantaneous phase vectors $\mathbf{u}_t = e^{i\boldsymbol{\theta}_t}$ retained. A model can then be fitted to a pool of phase vectors, explaining consistent phase differences between two or more signals. As illustrated in Fig. 2, modeling only the real part (cosine) of the complex phase vectors as suggested in, among others, [21,22], results in phase shifts of θ and $2\pi - \theta$, i.e., reflections around the real axis, becoming indistinguishable. Similarly, modeling the imaginary part of coherence [26] results in signals that are in phase becoming indistinguishable from signals in antiphase (reflection around the imaginary axis). Finally, only the complex representation can account for all consistent phase deviations as opposed to random, uncorrelated signals. Consequently, the full complex representation is needed to account for all possible types of phase coherence and previous methods focusing on, e.g., only cosine or sine coherence are inherently limited in their representational value.

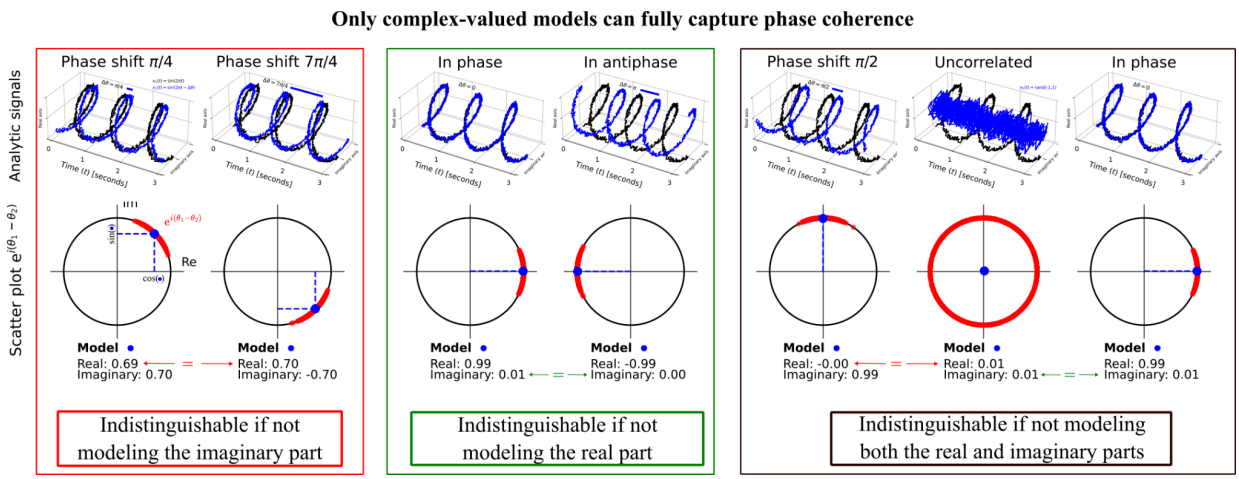

Figure 2: Modeling phase coherence in its entirety. Two oscillatory signals can be compared by examining their phases, which indicate the position of each sample within the oscillation cycle. The phase difference of such signals indicates their average displacement (phase shift). Gathering many phase time series, e.g., from fMRI data, enables modeling the coherence between all region pairs. However, previous models have resorted to only modeling the cosine of pairwise phase differences, meaning that phase shifts θ and $2\pi - \theta$ (red box) would lead to the same model estimate. Likewise, phase shifts of $\pm\pi/2$ cannot be distinguished from uncorrelated signals (black box) using cosine models. Similarly, modeling only the sine, corresponding to the imaginary part, fails to distinguish in-phase from antiphasic signals (green box) and in-phase from uncorrelated signals (black box). These issues are resolved when modeling both cosine and sine in a combined complex-valued representation.

To confirm the practical implications of not modeling the full complex-valued phase coherence, we generated synthetic data from $K = 2$ true components with consistent phase shifts and varying noise levels in $p = 3$ dimensions (Supporting Fig. S1). As expected, we found that mixture models for the complex-valued phase coherence, including the CACG model, best recovered the original data distribution.

Rank-1 component covariances are insufficient for modeling anisotropic data

Multivariate brain data is anisotropic over time; i.e., two adjacent brain areas are likely to have similar neuroimaging time-series while a third, unrelated brain area may have a time-series unrelated to the first two (Fig. 3A). In this case, the data distribution tends to a low-rank manifold (or subspace), e.g.,

a hyperplane in the Euclidean vector space or a great circle for phase data on the projective hyperplane. Importantly, the rank of the subspace can fall anywhere between one and the data's dimensionality.

To account for anisotropy in data while avoiding overparameterization, a covariance matrix parameter $\mathbf{\Psi}$ may be reparameterized using a low-rank matrix plus the identity matrix, i.e., the so-called low-rank-plus-diagonal reparameterization: $\mathbf{\Psi} \approx \mathbf{Z} = \mathbf{M}\mathbf{M}^{\mathsf{H}} + \mathbf{I}$. Here, $\mathbf{M} \in \mathbb{C}^{p \times r}$ is a rank-r matrix, $\mathbf{I}$ is the identity matrix, and $(\cdot)^{\mathsf{H}}$ is the complex conjugate transpose, reducing to the transpose for real-valued $\mathbf{M}$ [34]. This allows explicitly controlling the complexity of connectivity captured by the components of the mixture model. For some distributions, such as the Gaussian distribution, a scale-parameter on the identity matrix is also required (Methods).

A two-component mixture model with rank-controlled CACG distributions (real-valued for visualization purposes) as mixture components capture the data to differing degrees (Fig. 3B): Rank-1 components are inherently isotropic, for which potentially infinitely many mixture components are needed to account for the anisotropy of the data. Instead, rank-2 mixture components sufficiently explain the data. Additional components or increased component rank do not provide any gain in the predictive likelihood; such additional parameters are redundant and should be removed in keeping with the law of parsimony. Of note, many K-means models and matrix decompositions have rank-1 components and should thus be used with caution if the data is anisotropic.

Intercorrelations in fMRI data are seen by the naturally occurring network groupings of brain areas requiring modeling covariances beyond a rank-1 representation (i.e., non-unit rank) (Fig. 3C). Contrarily, a full covariance parameter $\mathbf{\Psi}$ can become singular leading to exploding normalization constants or modeling excessive numbers of parameters without explanatory gain.

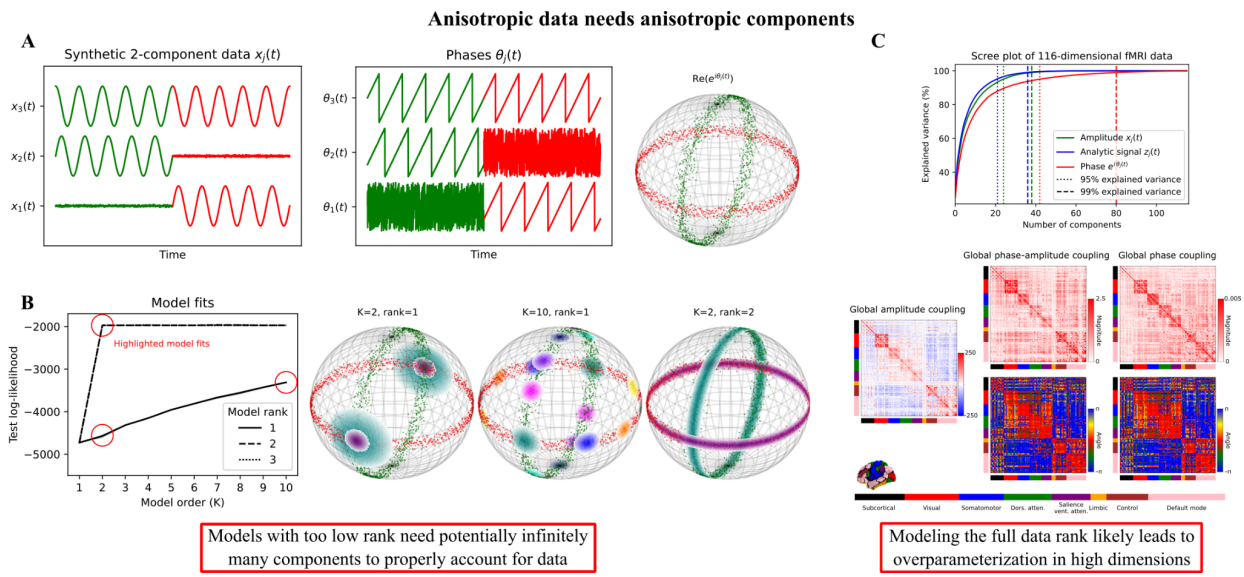

Figure 3: Non-unit rank models are needed to account for anisotropy present in functional neuroimaging data. ***A****): Real-valued synthetic anisotropic phase vectors in two synthetic components.* ***B****): The (real-valued) angular central Gaussian mixture model with rank-1 components fails to account for the data anisotropy and instead fits isotropic components along the data manifold. With rank-2 or rank-3 covariance, the model adequately covers the spherical subspace on which the data is distributed.* ***C****): Functional neuroimaging data is naturally grouped into networks, which are inadequately described by rank-1 components but likely overparameterized by full-rank models. The solution involves modeling a certain predefined rank r in between 1 and p. Coupling matrices shown are the amplitude covariance corresponding to the conventional covariance matrix of the time-series data $x(t)$, phase-amplitude coupling accounting for both phase and amplitude based on the covariance of $z(t)$, and phase coupling disregarding amplitude information based on the covariance of $u(t)$ shown as the global average covariances across the 255 subjects from the Human Connectome Project in resting state.*

To assess the practical implications of component ranks, we controlled the phase information in resting-state fMRI data from one subject by manipulating the phase in the Fourier domain into $K = 5$ segments (Supporting Fig. S2). Model performance was observed to stagnate around $r = 10$ with $p = 116$, indicating that unit-rank models underparameterize the data and full-rank models likely overparameterize. CACG mixtures outperformed all other models in recognizing the ground truth segments, showing their superiority in detecting phase changes in data unrelated to amplitude. Furthermore, we observed no performance benefits of using hidden Markov models over mixture models (Supporting Fig S3) and therefore opted to only present mixture model results for the rest of the study.

These mixtures were estimated using nonlinear optimization in PyTorch as we observed no benefit of analytical estimation via the expectation-maximization algorithm (Supporting Fig. S3).

Phase coherence mixture modeling provides meaningful partitions into task-related functional states

Next, we investigated whether a low-rank phase coupling mixture model can extract meaningful task-related patterns from HCP fMRI (emotion, gambling, language, motor, relational, social, and working memory). Each task scan included both active and resting periods but we did not distinguish between these blocks in our assessment; instead, all volumes in a task scan were labeled as part of the task. We applied the hold-out method and left 99 subjects out for testing (one subject excluded for not having all tasks completed) keeping 155 train subjects for training. We fitted seven-component mixture models with varying component rank to the pooled training data ($155 \cdot 1940$ fMRI volumes, no labels) and subsequently evaluated test-set posterior probability overlap with true task labels using normalized mutual information (NMI, Methods) as well as the volume-wise and scan-wise prediction accuracy of the seven classes. Each model was run 10 times; only the best log-likelihood fit was analyzed (Fig. 4).

The best performing ranks were non-full (i.e., less than p) for all models: CACG for phase coupling (best rank r=25), zero-mean Gaussian for amplitude coupling (r=100, minor gain over r=50), and zero-mean complex Gaussian for phase-amplitude coupling (r=50) (Fig. 4A). The CACG ($r = 25$) achieved a test-NMI of 0.44 (volume-wise mean accuracy 52%, chance=14%), significantly outperforming the optimal Gaussian (32%) and complex Gaussian (34%) fits (Wilcoxon signed-rank $p_{\text{bonf}} < 0.001$). Scan-wise (one predicted label per scan) accuracy was higher (CACG: 73% significantly outperforming Gaussian: 40%, complex Gaussian: 44%, $p_{\text{bonf}} < 0.001$). All mixture models performed better than rank-1 K-means baselines, of which LEiDA performed the worst.

To understand how model performance varies depending on the input data, we repeated the analysis at varying parcellation size and without global signal regression (GSR). CACG performance decreased without GSR (36% vs 52% with GSR), whereas amplitude models were less affected: 32% for Gaussian ($p_{\text{bonf}} < 0.001$ vs CACG) and 38% for complex Gaussian ($p_{\text{bonf}} = 0.25$ vs CACG, Supporting Fig. S4). At higher parcellation size, low-rank models performed better (rank-25 CACG accuracy at 232 regions: 57%) but high-rank models performed worse, indicating that controlling the number of estimated parameters becomes increasingly important with higher dimensionality. CACG was sensitive to an alternative denoising with temporal ICA, which reduced CACG accuracy to 44% while Gaussian (32%)

and complex Gaussian (37%) were unaltered and still significantly worse than CACG ($p_{\mathrm{bonf}} < 0.001$), suggesting GSR and preprocessing influence recoverable phase structure.

The average test set posterior indicated that the language, motor, and social tasks were most distinct, while emotion, relational, and working memory tasks were less discernible (Fig. 4D); gambling was the hardest to detect, likely due to the very short event design (1s blocks) leading this task to be excluded in other task classification studies on HCP data [35,36]. One mixture component (Component 1) behaved like a background state, indicating the presence of a non-task-related background component with spurious activation. The remaining six components were well-mixed across tasks and subjects (component weights ranged between 0.12–0.18). The posterior also indicated consistent within-task fluctuations, e.g., five peaks for the social task, which is composed of five task blocks, and three peaks for the relational task, which is composed of three blocks. Confusions in CACG were mainly driven by gambling, relational, and working memory (Fig. 4E). The posteriors for the Gaussian and complex Gaussian showed lower task discernibility (Supporting Figs S5-6); these models were skewed in favor of a subset of components with high weight, leading to low-weight components specific to (particularly) high-amplitude observations even if these are of low quantity.

Each component was characterized by a covariance matrix which we summarize in 8 canonical networks [37]. Few between-network connections were consistent across all components (Fig. 4F): the salience network was always coherent with somatomotor and dorsal attention networks, and default mode with limbic and control. The default network was always anticoherent with the dorsal attention network, visual with limbic, somatomotor with control, and salience with limbic. We summarize each CACG component by network connections that were significantly characteristic of that component versus all other components in terms of subject-wise connection strength (max p-value across component pairs using paired Hotelling's T2-test [38] on test subjects, Bonferroni corrected). The subcortical network only had two significant connections across components while the visual network had 17, the remaining networks had between 7-12 significant connections indicating brain-wide task engagement.

Component 1 (low task engagement, blue) showed generally low overall coupling with the notable exception of within-network visual coherence likely indicating the eyes-open condition. Dorsal attention–limbic coupling was significantly reduced ($p_{\mathrm{bonf}} < 0.05$), pointing to this connection serving as a potential task-activation pathway. Component 2 (emotion, yellow) reflected a structured, arousal-like brain-wide state with broadly increased connection strength: limbic and default networks displayed stronger within-network coherence and visual-dorsal attention and limbic-default coherences were

increased, similarly with dorsal attention-default, dorsal attention-limbic, visual-default, visual-limbic, and salience-limbic anticoherence. Component 3 (language, green) was characterized by increased control-network coherence internally and pronounced anticoherence between control and subcortical, visual, and somatomotor networks; unlike most components, visual and somatomotor networks were coherent, visual–dorsal attention showed anticoherence, whereas the somatomotor–salience coherence was reduced. The control network includes some language areas (left inferior frontal gyrus, angular and supramarginal gyri), so activation of this network is not surprising. Component 4 (motor, red) exhibited significantly reduced somatomotor within-network coherence and weakened anticoherence with visual regions; dorsal attention was coherent with somatomotor and consistently phase-shifted relative to visual, while visual was phase-shifted relative to default (a pattern unique among components), and salience–default anticoherence was strengthened. This finding suggests that area-specific motor engagement does not lead to network-wide somatomotor coherence. Component 5 (relational, purple) was marked by strengthened anticoherence between visual–default, visual–limbic, visual–somatomotor, and somatomotor–dorsal attention, alongside coherent visual–dorsal attention and somatomotor–default coupling (distinct from other components); portions of somatomotor–subcortical interactions were consistently phase-shifted. This component could indicate a perceptual integration mode, wherein visual information is mapped onto higher-order representations. Component 6 (social, brown) showed significantly increased control within-network coherence and enhanced visual–dorsal attention coherence, while control–dorsal attention relationships were anticoherent and visual–salience interactions were phase-shifted relative to most other components. The control-visual-attention coupling could indicate a structured reading of social gestures in the task. Finally, Component 7 (working memory, pink) was defined by within-network coherence in visual and control systems together with anticoherence between somatomotor and control networks, indicating that the control network collaborates with the visual network to actively memorize visual information.

Taken together, the CACG low-rank mixture model partitions the data into components that align with task structure and reveal uniquely interpretable changes in network coherence strength and phase polarity.

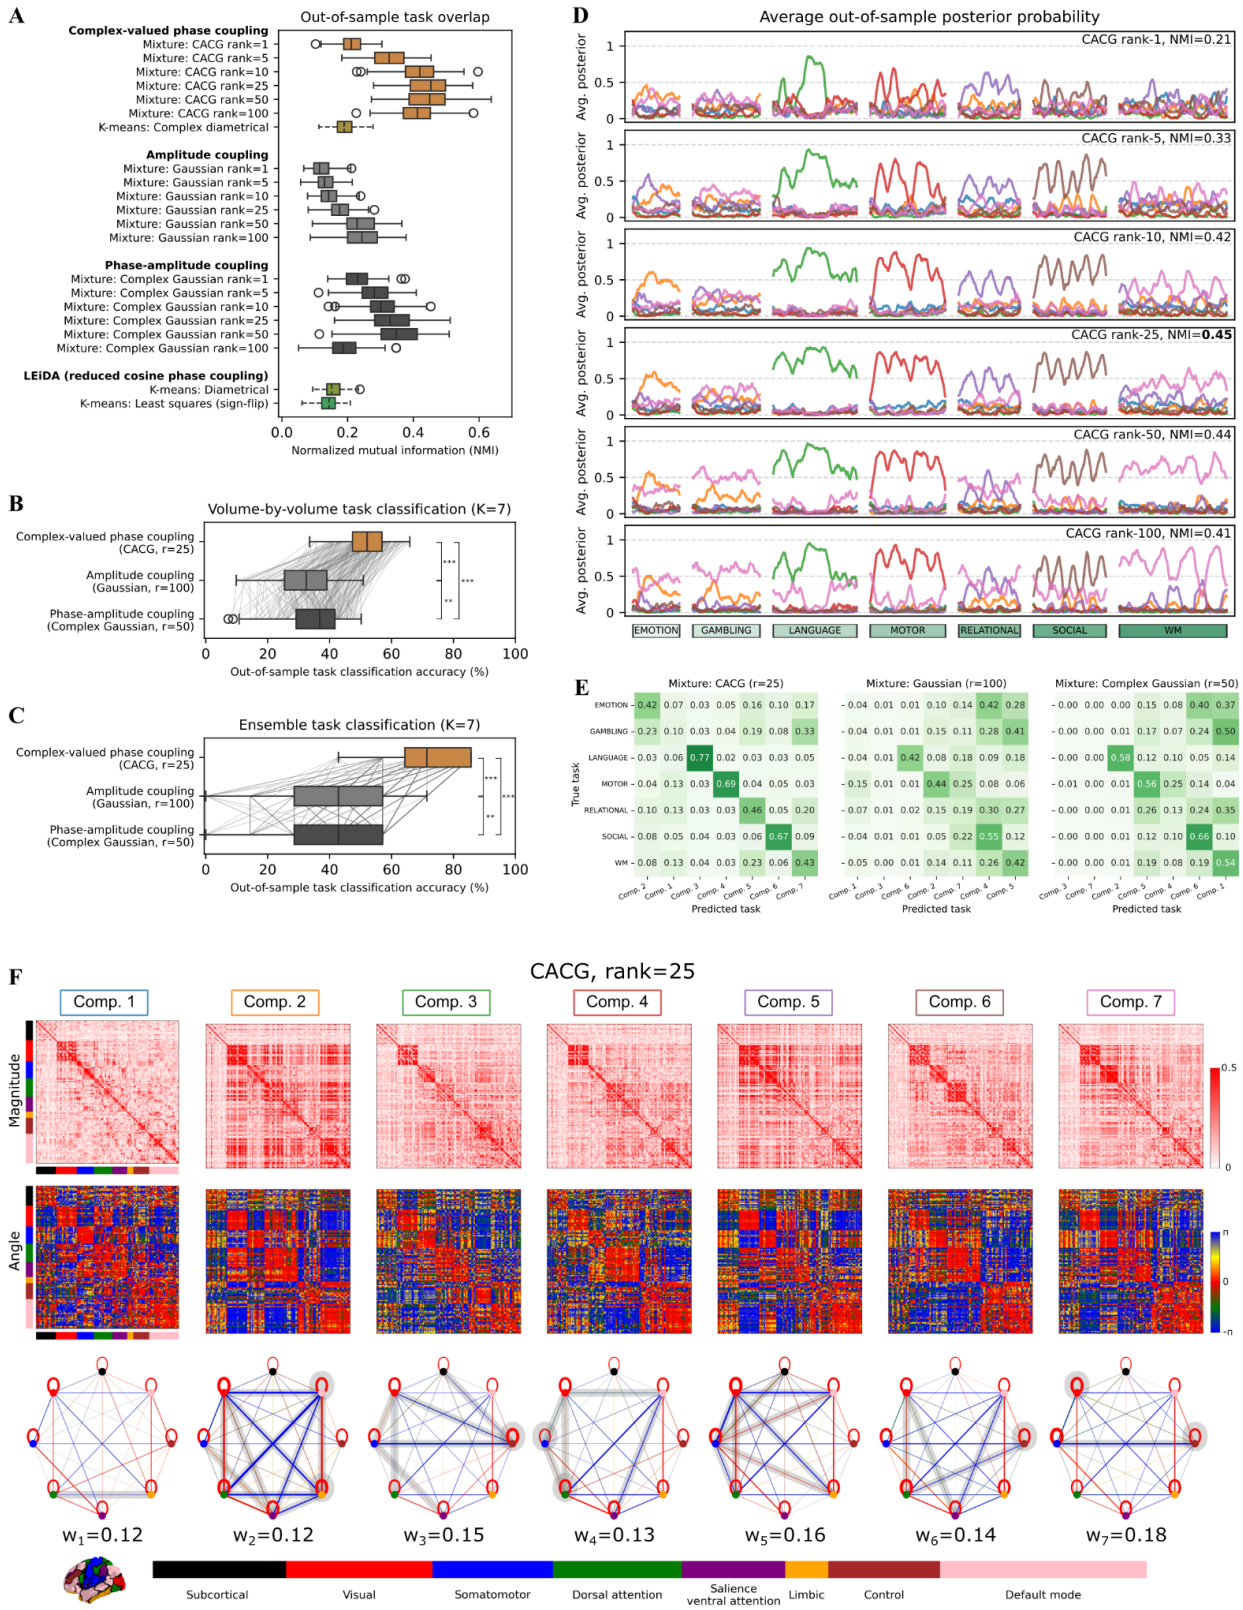

*Figure 4: Seven-component mixture models applied to the seven task fMRI scans from the human connectome project (HCP). **A**): Overlap between posterior probability and the true task labels measured*

using normalized mutual information. ***B****): Volume-by-volume task classification accuracy using binarized test set posterior probabilities. Components were matched to tasks based on their average performance across train subjects.* ***C****): Same as B), with only one predicted component per scan. In A), B), and C), boxplots depict values for N=99 test subjects. In B) and C), ** $p < 0.01$, *** $p < 0.001$ indicates significance in a Wilcoxon signed rank test, Bonferroni-corrected for three comparisons.* ***D****): The average test posterior of each of the seven components (one color for each) of the CACG mixture model at different component ranks. The x-axis represents time with whitespace demarcating different task scans. Boxes indicate the name of the task and box widths are proportional to the number of volumes in the task scan.* ***E****): Test-set confusion matrices after ordering components to their most likely task as measured by train accuracy.* ***F****): The seven components learned using the best-performing model, i.e., the CACG mixture model with rank-25 components. In the network depiction, the color of edges corresponds to the average interregional "angle" and the line thickness to the average interregional magnitude in 8 canonical networks. Significant connections (Hotelling's T2-test, $p_{\text{bonf}} < 0.05$) are highlighted with a grey background.*

Resting-state phase coherence is better explained by higher component rank than more components

For the resting-state data we evaluated the number of mixture model components and component rank using three subject-specific metrics on the test data set: 1) predictive log-likelihood indicating how well the test data matched the learned mixture distribution, 2) entropy of the average posterior indicating how "mixed" the components were for each subject, and 3) computational reproducibility indicating how sensitive the estimate was to changes in initialization influenced by the presence of local minimas.

For the CACG mixture model, increasing the number of components altered likelihoods less than increasing component rank (Fig. 5A). When model rank tended towards the data's dimensionality ($p = 116$), fewer components were supported by the test-set resting state data, and three components with rank-100 components achieved the highest predictive likelihood. These results suggest that resting-state brain fMRI data is not easily split into temporal segments that are representative across subjects and scans. Within-subject component mixing was slightly lower for ranks 50 and 100, indicating that the model tends to segregate subjects rather than individual volumes when components are of high rank. Computational reproducibility was highest for $K \leq 4$ across all component ranks ($NMI > 0.8$ for most models). For higher model orders, estimates tended to fluctuate more, indicating that these models

benefit from repeated initialization. Conclusions were similar for the Gaussian and complex Gaussian mixture models (Supporting Figs S8-9).

We present the three-component CACG solution with rank-25 components, which had high within-subject mixing and computational reproducibility in Fig. 5B. All components largely grouped sensorimotor and attention networks together against default-limbic and default-control networks. Given the fewer components, more connections were significant across components than for the task analysis; here we only describe selected connections. Component 1 was characterized by strengthened dorsal attention and somatomotor within- and between-network coherence and anticoherence between somatomotor and control networks, indicating task-positive activation potentially related to movements. Component 2, which was the most prominent, was characterized by decreased somatomotor/attention and control connectivity but otherwise displayed the conventional dissociation between default and task-positive networks. Component 3 was characterized by reduced dorsal attention and salience connectivity. The estimated components were similar to those estimated using Gaussian and complex Gaussian components (Supporting Figs. S7-8) though these seemingly had a wider distribution of network connectivity strengths.

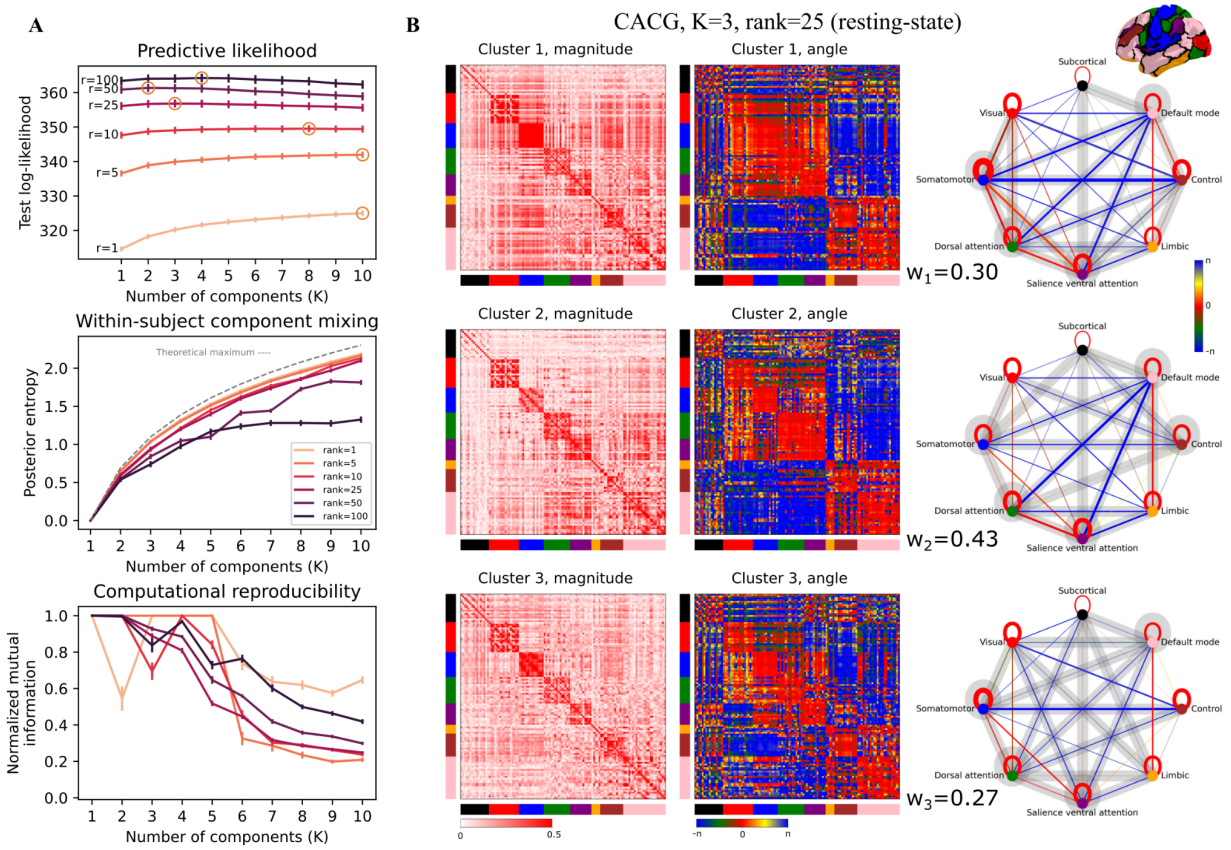

Figure 5: Mixture model parameter search and estimate on resting-state fMRI for the CACG model. ***A****): Component rank and model order compared with out-of-sample test log-likelihood (normalized by the number of samples), posterior entropy indicating component mixing for each subject, and computational reproducibility upon repeated initializations of the same model; bars represent standard error across N=99 test subjects.* ***B****): CACG mixture estimate with K=3 and component ranks* $r = 25$*. In the network representation, average inter-network angle is represented by color, and magnitude by edge thickness. Significant connections (Hotelling's T2-test,* $p_{\mathrm{bonf}} < 0.05$*) are highlighted with a grey background.*

GSR significantly altered the average resting-state global phase coherence map, revealing interregional anti-coherence in the angular representation (Supporting Fig. S9). Without GSR, the vast majority of the brain was, on average in this data, coherent. Following GSR, the average angular global phase coherence was predominantly close to zero (in-phase) or π (anti-phase). Visually, a CACG model with rank 10 was able to fit the global average phase coherence map well, with a visually negligible effect of increasing the rank. In contrast, a unit-rank model provided a notably simplified estimate of the global coherence map, essentially dividing the cortical areas into two chunks, again suggesting that rank-1 models are too restrictive.

DISCUSSION

Here we explored and advanced multivariate probabilistic mixture modeling approaches for analyzing large-scale networks of phase coherence in their entirety, emphasizing the importance of discarding as little potentially valuable information as possible when modeling the data. Based on our results, we argue that simplifying data and models—such as by discarding the imaginary component of phase coherence, relying on leading eigenvectors, or utilizing rank-1 components—can result in loss of cognitively relevant information. Our findings show that appropriate statistical distributions that focus on phase coherence rather than signal amplitude better retrieve task information and provide a clean representation of brain dynamics. Importantly, our models were trained on label-free data and evaluated on unseen data from multiple subjects, indicating that the estimated models provide generalizable representations of dynamic brain connectivity.

We introduced mixture models that were entirely agnostic to task labels, yet nonetheless supported accurate classification in held-out subjects (57% accuracy for the phase-synchronization model in 232 parcels). Although task decoding was not a primary objective, this performance indicates that the learned

unsupervised models capture cross-subject brain dynamics well enough to distinguish cognitively meaningful brain states. Prior work on task classification in the HCP dataset has relied largely on supervised or deep learning approaches. Classical supervised models have typically been applied to static functional connectomes [39] or restricted to the immediate post-stimulus interval [36], limiting their ability to provide per-volume interpretability. Supervised discriminative methods also optimize directly for task labels and often depend on summary features, which further constrains mechanistic insight. Deep learning approaches achieve near-perfect decoding accuracy [40,41], but this is generally accompanied by reduced explainability of the underlying neural processes driving the predictions. In contrast, our framework offers a complementary and mechanistically interpretable account of dynamic brain organization. It preserves sample-level temporal resolution, provides meaningful task-related structure without requiring labels, and thereby bridges the gap between decoding performance and interpretability. By providing a principled way to extract and compare latent components across task and resting-state data, our approach enables future investigations to test specific hypotheses about behavioral relevance, physiological correlates, and underlying neural mechanisms. Applying this framework in conjunction with experimental perturbations or richer behavioral measures represents a natural next step.

Through synthetic data, we illustrated that focusing solely on cosine phase coherence may mask valuable information about phase shifts between signals, a drawback remedied by complex-valued components (Fig. 2). The estimated seven-component CACG model for task fMRI data also suggested some valuable information in the imaginary component (green/yellow edges in Fig. 4F, “angle”). In the synthetic, phase-controlled and task-fMRI experiments, complex-valued models performed better than LEiDA, which models the leading eigenvector of the cosine phase coherence matrix and has been used extensively in the literature (Figs. 4, S1-S2, S4). Taken together, constraining phase coherence models to the real domain risks discarding potential insights embedded in the imaginary component without obvious benefit.

In multivariate statistical modeling, attention to the appropriate data manifold is crucial. Eigenvectors, which are by definition unit norm and sign-invariant, lie on the projective hyperplane, also known as the sign-symmetric unit hypersphere. Assuming the wrong manifold, such as using least-squares or cosine distance methods on eigenvectors, can introduce inaccuracies, even with eigenvector sign-flipping as performed in LEiDA, which imposes non-affine restrictions on the data space [42]. Some studies have suggested a K-means approach using either the l1-norm, cosine distance, or Euclidean distance between the upper-triangular part of the full cosine phase coherence maps and centroids of the same size [43–45], neither of which respect the manifold distribution of phase data. Furthermore, representing the triangular

part of a symmetric matrix disregards its Riemannian structure leading to systematic modeling errors. Instead, models that align with the correct manifold and allow for components of variable rank are more suitable.

We estimated our models using both analytically derived update rules via the expectation maximization (EM) as well as numerical optimization in PyTorch. We generally found the two methods to perform similarly (Supporting Fig. S3) with the latter being faster for large data sets. However, since the EM algorithm is more consistent (guarantees monotone convergence in many cases), we supply update rules and suggestions for model optimization for the CACG distribution based on EM, including low-rank-plus-diagonal reparameterization, in the Supporting Information. While our analyses were limited to low-dimensional data, the low-rank-plus-diagonal reparameterization combined with gradual initialization of high-rank models (Methods) support the feasibility of future models in high, potentially native, resolution.

We showed that the global average phase coherence matrix without GSR is highly homogeneous (Supporting Fig. S9). Clustering cosine phase coherence maps without GSR typically results in one dominant global component with a high weight [46,47], while other clusters, previously referred to as ghost attractors [48], carry less weight. The CACG distribution for phase data was markedly more affected by the removal of GSR than amplitude and phase-amplitude models (Supporting Fig. S4). Without GSR, fMRI data is often dominated by the global signal leading data to conform largely to this rank-1 component. Without access to amplitude information, the mixture of CACG distributions struggles more to discern subcomponents within such data. This is supported by high-rank CACG models, specifically, performing worse without GSR. In Gaussian and complex Gaussian models, amplitude variation allows for differentiating samples conforming to the global signal. Thus, for phase data, GSR is not just a denoising step to remove potential confounds, it is a necessity to extract meaningful phase coherence components. We believe that applying GSR, which removes more trivial coherence structures, will directly lead to learned components having more contrast, characterizing more detailed and relevant connectivity information.

A disadvantage of synchronization methods employing the Hilbert transform is the requirement of data to be narrowband for phase estimation. We used a filter in the range $[0.03, 0.07]Hz$ based on previous studies [21,23,24,49], which is a relatively small restriction on the $[0.01, 0.1]Hz$ range that the BOLD signal is generally accepted to fall within. Future studies should further investigate the effect of bandwidth on phase estimation in BOLD fMRI data, e.g., focusing on detecting phase slips [50]. Previously,

the empirical mode decomposition has been suggested as a data-driven method to extract narrowband signal "modes" prior to the Hilbert transform [51]. However, it is difficult to reconcile frequency bands of interest across subjects and in practice, this method subdivides the frequency axis log-linearly with little indication of the relative relevance of each derived frequency band.

Besides the frequency range, our study has other limitations. The computational training speed increased markedly from K-means (time scale of seconds for 155 subjects) to mixture modeling (minutes) and hidden Markov Modeling (hours). However, we did not exploit GPU-processing, which is readily available in PyTorch and could be implemented in future studies. We compared methods for phase coherence with zero-mean Gaussian mixture models for "raw" time-series data, however, in the literature, it is more common to see hidden Markov models for fMRI data, perhaps augmented with autoregressive parameters [19,52–55]. Here we observed little difference between mixture and hidden Markov modeling (Supporting Fig. S3). With the PyTorch stochastic estimation framework, future studies could very easily add an autoregressive emission specification to the presented models. Finally, we only examined one dataset; the estimated components may very well differ in different data.

Conclusion

Modeling time-varying spatiotemporal coherence of functional brain imaging signals can provide meaningful insights about neural systems underlying cognition and illness. Here we demonstrated that dynamic macroscale functional brain phase coherence networks can be modeled in their entirety considering mixture modeling approaches with 1) complex-valued and 2) non-unit rank centroids. We demonstrated the efficacy of the CACG mixture model on both synthetic and real resting-state and task fMRI. We suggest that future studies within phase coherence, whether in neuroscience or other fields, consider models such as the CACG distribution that account for anisotropy and phase coherence in its entirety. Our Phase Coherence Mixture Modeling (PCMM) toolbox is readily available for general-purpose phase coherence modeling: github.com/anders-s-olsen/PCMM.

METHODS

Table 1: Signal preprocessing considerations regarding phase estimation via the Hilbert transformation. The processing steps are outlined for a univariate signal and should be repeated across ROIs/voxels $j = 1, \ldots, p$. *Padding size in step one should be large enough to reasonably cover the Hilbert impulse response and be applied to both the front and back end of the signal in the case of zero-phase (forward-reverse) filtering. Scipy.signal.filtfilt() pads by default but scipy.signal.hilbert() does not, which is why we suggest a consistent padding throughout the pipeline.*

Signal preprocessing step	**Reason**	**Example numpy/scipy code**
Pad signal $x_j(t)$ with mirrors of itself.	To avoid edge effects from bandpass filtering and the Hilbert transform (the effect is smaller for phase than for amplitude data).	`pad = 100` `x_j_padded = numpy.concatenate([x_j[1:pad+1][::-1], x_j, x[-pad-1:-1][::-1]])`
Narrowband bandpass filter, e.g., $0.03 - 0.07Hz$ for BOLD fMRI.	Phase estimation requires "narrowband" data. The best signal bandwidth is unclear.	`b, a = scipy.signal.butter(N=2, Wn=[0.03,0.07], btype="bandpass", fs=fs)` `x_j_filtered = scipy.signal.filtfilt(b, a, x_j_padded, padtype=None)`[56]
Subtract the mean.	Phase estimation introduces alterations for signal mean above zero and breaks down for signal mean above one [56].	`x_j_filtered_demeaned = x_j_filtered - numpy.mean(x_j_filtered)`

Decompose signal into phase $\theta_j(t)$ and magnitude $a_j(t)$ using the Hilbert transform.	Separates the oscillatory and amplitude constituents of the signal.	`x_j_hilbert = scipy.signal.hilbert(x_j_filtered_demeaned)`
		`a_j = numpy.abs(x_j_hilbert)`
		`theta_j = numpy.angle(x_j_hilbert)`

Trim the obtained magnitude and phase signals to the original length.	To avoid artificially increasing the signal length via padding.	`a_j = a_j[pad:-pad]` `theta_j = theta_j[pad:-pad]`
Convert phase from polar to rectangular format and rescale by root dimensionality .	Enables modeling of phase data on the complex projective hyperplane using, e.g., the complex ACG distribution.	`u_j = 1/numpy.sqrt(p)*numpy.exp(1j*theta_j)`

Optional: Save both magnitude and phase for modeling or only the phase.	Modeling only the phase $\{u_j = \frac{1}{\sqrt{p}} e^{i\theta_j}\}_{j=1}^{p}$ ensures an exclusive focus on *synchronization* and avoids amplitude-specific artifacts.	`data[:,j] = u_j # complex phase vector`
		`data[:,j] = a_j * u_j # analytic signal`
		`data[:,j] = x_j_filtered_demeaned # filtered time-series`

Experimental data and preprocessing

We used fMRI data from 255 unrelated healthy subjects from the openly available Human Connectome Project [31,57]. Ethics statements and imaging protocols may be found elsewhere [58,59]; we downloaded the minimally preprocessed, multimodal surface matched task and resting-state data, the

latter had also been ICA-FIX denoised [60]. We only used right-left phase encoding scans. Of the 255 subjects, 155 randomly selected individuals were used for training, while the remaining 100 subjects were exclusively used as an out-of-sample test set. For the 155 subjects, only one of two scans (REST1 or REST2, 1200 samples each, randomly selected) were included in the training set. One subject was discarded from the task analysis out-of-sample test set due to not having all task scans complete. We also evaluated a 2025-version of the HCP dataset using updated temporal and spatial ICA artifact classification schemes [61]. The task scans consisted of emotion (176 volumes), gambling (253), language (316), motor (284), relational (232), social (274), and working memory (405) tasks with a total of 1940 volumes per subject, see [58] for a full description of tasks.

The data was linearly detrended voxel-wise and subsequently spatially downsampled to 100 cortical and 16 subcortical regions using the Schaefer and Tian atlases [32,33]. Subsequently, we bandpass-filtered the data in the range $[0.03, 0.07]Hz$ with a 2nd-order zero-phase butterworth filter. Due to the relatively short acquisition lengths and long impulse response of the filter giving rise to edge effects from both filtering and the Hilbert transform, we prepended and appended each signal with 100 flipped samples of the signal itself, which corresponds exactly to 1% of the peak of the filter impulse response for TR=0.72s. The added samples were discarded after the Hilbert transform. Importantly, phase estimation via the Hilbert transform requires narrowband data; however, no direct guidelines for the required signal bandwidth exist, only rules of thumb such as requiring the bandwidth to be smaller than the mid-frequency [62] or twice the mid-frequency [63], both of which we enforce using $0.03 - 0.07Hz$ as our frequency range. See Table 1 for a summary of the key processing steps required for phase coherence modeling.

Dynamic interregional phase coherence

Whereas the Pearson correlation coefficient between regional neuroimaging time-series implicitly assumes that the brain covariance matrix is static, time-varying phase coherence can be estimated via the Hilbert transform. Evaluating dynamic brain coherence implicitly assumes that the signal in all brain areas is of the same scale and that coactivity between regional time courses can be wholly represented by the coherence between their phases.

The regional neuroimaging signal $x_j(t)$ for region $j \in \{1, \ldots, p\}$ is extended into the complex domain using the Hilbert transform $x_j^{(h)}(t) = x_j(t) * \frac{1}{\pi t}$, where $*$ is the convolution operator, to form the analytic signal

$$z_j(t) = x_j(t) + i x_j^{(h)}(t).$$

The Hilbert transform represents a phase shift by $\pi/2$ radians to every frequency in the signal. While the original signal $x_j(t)$ is oscillating along the time axis, the analytic signal $z_j(t)$ is revolving around the unit circle in the complex plane (Figs. 1-2). In the complex plane every point may be described using polar coordinates, i.e., as a combination of an angle to the positive real axis

$$\theta_j(t) = \tan^{-1}\left(\frac{x_j^{(h)}(t)}{x_j(t)}\right),$$

also called the instantaneous phase, and a distance from the origin $a_j(t) = \sqrt{x_j(t)^2 + \left(x_j^{(h)}(t)\right)^2}$, known as the instantaneous magnitude. The phase captures solely the oscillatory information in the signal while signal strength, including potential artifactual spikes, forms part of the signal amplitude. Whereas the Hilbert transform itself does not impose frequency restrictions on the signal, reliable estimation of phase requires that the signal be sufficiently narrowband [64]. Furthermore, while the naming convention "instantaneous" implies that phases can be calculated using information from only a single time point, in reality the Hilbert transform suffers from the same estimation requirements for a sufficient number of samples and linearity/stationarity assumptions as the Fourier transform.

Each sample of the analytic signal is a complex number $z_j(t) = a_j(t)e^{i\theta_j(t)}$, where the radius is ignored when exclusively considering phase coherence, i.e., $a = 1$. The phase θ and difference between two phases $\theta_1 - \theta_2$ are circular variables with modus 2π. The rectangular representation can therefore be used to assess coherence via $e^{i(\theta_1 - \theta_2)}$ and more generally, all pairwise complex instantaneous phase coherences can be represented by the matrix $\mathbf{A}_t$ with elements

$$A_{t,j,j'} = e^{i(\theta_{t,j} - \theta_{t,j'})} = \cos(\theta_{t,j} - \theta_{t,j'}) + i\sin(\theta_{t,j} - \theta_{t,j'}),$$

such that every element for nodes $j, j' \in \{1, \ldots, p\}$ provides an estimate of the pairwise instantaneous synchrony between regions. Each element of $\mathbf{A}_t$ can be factored as $A_{t,j,j'} = e^{i(\theta_{t,j} - \theta_{t,j'})} = e^{i\theta_{t,j}} e^{-i\theta_{t,j'}}$, and thus, the matrix can be written as $\mathbf{A}_t = \mathbf{u}_t \mathbf{u}_t^{\mathsf{H}}$, where $\mathbf{u}_t = e^{i\boldsymbol{\theta}_t}$ and $(\cdot)^{\mathsf{H}}$ represents the complex conjugate transpose. Thus, the rank of $\mathbf{A}_t$ is 1, and instantaneous interregional phase coherence can be represented using a vector, $\mathbf{u}_t$, instead of a matrix with $\frac{p(p-1)}{2}$ unique elements, greatly simplifying the

modeling of all pairwise coherences. Since $\mathrm{diag}(\mathbf{A}_t) = \mathbf{1}$, i.e., a vector of all ones, the only non-zero eigenvalue of $\mathbf{A}_t$ is $\lambda = p$. Thus, $\mathbf{u}_t$ may readily be rescaled to unit norm, $\mathbf{u}_t = \frac{1}{\sqrt{p}} e^{i\boldsymbol{\theta}_t}$. Combined with the fact that flipping the sign of $\mathbf{u}_t$ does not change $\mathbf{A}_t$, it means that $\mathbf{u}_t$ is an eigenvector and models of interregional phase coherence should therefore also be sign-invariant. Consequently, all pairwise instantaneous phases can be compactly represented using a complex-valued sign-invariant vector residing on the unit sphere, also known as the complex projective hyperplane $\mathbb{CP}^{p-1} = \{\mathbf{u} | \mathbf{u} \in \mathbb{C}^p, \mathbf{u}^{\mathsf{H}}\mathbf{u} = 1, -\mathbf{u} = \mathbf{u}\}$.

It is also possible to model instantaneous phase coherence in the real domain by discarding the imaginary part of the complex phase coherence. Especially in fMRI research, the cosine phase coherence matrix has been widely studied, particularly in the context of the Leading Eigenvector Dynamics Analysis (LEiDA) [22] considering the real symmetric matrix $\mathbf{A}_t^{(r)}$ with elements

$$A_{t,j,j'}^{(r)} = \cos(\theta_{t,j} - \theta_{t,j'}) = \mathrm{Re}(A_{t,j,j'}).$$

Notably, a matrix of cosine differences is rank-2 due to the angle difference identity

$$\mathbf{A}_t^{(r)} = \mathbf{c}_t\mathbf{c}_t^{\top} + \mathbf{s}_t\mathbf{s}_t^{\top},$$

where $\mathbf{c}_t = \cos\boldsymbol{\theta}_t = \mathrm{Re}(\mathbf{u}_t)$ and $\mathbf{s}_t = \sin\boldsymbol{\theta}_t = \mathrm{Im}(\mathbf{u}_t)$, which we have shown previously in [47]. Consequently, $\mathbf{A}_t^{(r)}$ can be expressed in terms of *two* eigenvectors with non-zero eigenvalue. The leading eigenvector of the cosine phase coherence map automatically accounts for at least 50% of the variance and empirically captures anywhere between 50% and 100% [44]. Several previous studies have focused on the leading eigenvector $\mathbf{u}_{t,1}$ of $\mathbf{A}_t^{(r)}$ in the LEiDA-framework [22]. These eigenvectors have traditionally been sign-flipped, such that the majority of elements are of the same sign, and subsequently clustered using least-squares K-means [22,46,48,65]. This approach is a heuristic, which we previously have shown may lead to inaccurate results (see Fig. S1 in [47]). In our previous work, we therefore suggested using diametrical clustering as a projective hyperplane alternative for LEiDA [47].

Statistical distributions and low-rank reparameterization

Below we present statistical distributions for analyzing multivariate brain data from amplitude, phase, and phase-amplitude perspectives; the presented methods and their characteristics are summarized in Supplementary Table S1. Further information on the CACG distribution may be found in the Supplementary Information.

Complex angular central Gaussian distribution

The complex angular central Gaussian (CACG) distribution for data on the complex projective hyperplane has density [27,28]

$$f_{\mathrm{CACG}}(\mathbf{u}|\mathbf{\Psi}) = \frac{\Gamma(p)}{2\pi^p|\mathbf{\Psi}|}\left(\mathbf{u}^{\mathsf{H}}\mathbf{\Psi}^{-1}\mathbf{u}\right)^{-p},\ \mathbf{u} \in \mathbb{CP}^{p-1}.$$

Here, $\mathbf{\Psi} \in \mathbb{C}^{p\times p}$ is a Hermitian-symmetric positive definite matrix identifiable up to a positive scale factor ($\mathbf{\Psi}$ is scaled such that $\mathrm{tr}(\mathbf{\Psi}) = p$ for convenience) and $|\cdot|$ is the matrix determinant. The CACG distribution is named based on its sampling procedure: samples from a centered multivariate complex normal distribution normalized to the sphere follows the CACG, i.e., $\mathbf{x} \sim N(\mathbf{0}, \mathbf{\Psi}),\quad \frac{\mathbf{x}}{||\mathbf{x}||} \sim \mathrm{CACG}(\mathbf{\Psi})$ [66]. The Bingham distribution, which serves a similar purpose, includes a more complicated normalization constant (the confluent hypergeometric function of matrix argument), which becomes intractable for non-trivial dimensionality, rendering the CACG distribution an attractive alternative. Maximum likelihood estimation is deferred to the Supplementary Information.

Gaussian distribution

The real-valued zero-mean multivariate Gaussian distribution has the density

$$f_{\mathrm{ZMG}}(\mathbf{x}|\mathbf{\Psi}) = \frac{1}{(2\pi)^{p/2}|\mathbf{\Psi}|^{1/2}}e^{-\frac{1}{2}\mathbf{x}^\top\mathbf{\Psi}^{-1}\mathbf{x}},\quad \mathbf{x} \in \mathbb{R}^p.$$

Similarly, the zero-mean complex-valued multivariate Gaussian distribution has the density

$$f_{\mathrm{ZMCG}}(\mathbf{x}|\mathbf{\Psi}) = \frac{1}{\pi^p|\mathbf{\Psi}|}e^{-\mathbf{x}^\top\mathbf{\Psi}^{-1}\mathbf{x}},\quad \mathbf{x} \in \mathbb{C}^p.$$

Low-rank-plus-diagonal reparameterization of $\mathbf{\Psi}$

Estimating covariance matrix parameters can be challenging due to the near-quadratic increase in parameters with dimensionality. Moreover, high-resolution brain connectivity networks exhibit rank deficiency due to related areas having highly correlated time-series. To address this, a low-rank-plus-diagonal reparameterization $\mathbf{\Psi} \approx \gamma\left(\mathbf{M}\mathbf{M}^{\mathsf{H}} + \mathbf{I}\right)$, with $\mathbf{M} \in \mathbb{C}^{p\times r}$ (or $\mathbf{M} \in \mathbb{R}^{p\times r}$ for the Gaussian), allows capturing a pre-specified rank $r \leq p$ of the learned parameter matrix [34]. The constant $\gamma \in \mathbb{R}$ for the CACG distribution corresponds to a trace-normalization, i.e., $\gamma = \frac{p}{\mathrm{tr}\left(\mathbf{M}\mathbf{M}^{\mathsf{H}} + \mathbf{I}\right)} = \frac{1}{o/p + 1}$ where $o = ||\mathbf{M}||_F^2$, which does not affect mixture model results due to the scale-invariance of the CACG distribution. For the zero-mean Gaussian and complex Gaussian

distributions, that are not scale-invariant, we used the reparameterization $\mathbf{\Psi} \approx \mathbf{M}\mathbf{M}^{\mathsf{H}} + \gamma\mathbf{I}$ such that the scale parameter affects the relationship between the added identity matrix and $\mathbf{M}$. The use of low-rank-plus-diagonal reparameterization required a re-derivation of existing analytical maximization updates to be with respect to $\mathbf{M}$ as detailed in the Supplementary Information.

Probabilistic mixture modeling

Mixture modeling involves computing centroids for K components, each corresponding to a unique probability distribution mixed with component weights $\mathbf{w} = \{w_k\}_{k=1}^{K}, \sum_k w_k = 1$. The mixture density for each observation $\mathbf{y}_t$, which may correspond to, e.g., a (complex) eigenvector $\mathbf{y}_t = \mathbf{u}_t$ (phase coupling, CACG), a scaled phase vector $\mathbf{y}_t = \mathbf{a}_t \odot \mathbf{u}_t$ (phase-amplitude coupling, complex Gaussian), the time-series $\mathbf{y}_t = \mathbf{x}_t$ (amplitude coupling, Gaussian), is defined as

$$f_{\mathrm{MM}}(\mathbf{y}_t|\mathbf{\Psi}, \mathbf{w}) = \sum_{k=1}^{K} w_k f(\mathbf{y}_t|\mathbf{\Psi}_k),$$

where $\mathbf{\Psi}_k$ is the covariance matrix of the kth component.

We utilized two estimation schemes for mixture models involving these distributions: expectation-maximization (EM) and direct likelihood optimization. Both methods optimize the mixture model log-likelihood function evaluated over the entire data set $\mathcal{Y} = \{\mathbf{y}_t\}_{t=1}^{N}$

$$\mathcal{L}_{\mathrm{MM}}\left(\mathbf{\Psi}, \mathbf{w}|\mathcal{Y}\right) = \sum_{t=1}^{N} \ln\left(\sum_{k=1}^{K} w_k f(\mathbf{y}_t|\mathbf{\Psi}_k)\right).$$

EM computes assignment probabilities (i.e., state responsibilities a.k.a. posterior) β_{kt} for sample t and component k in the expectation (E-step) according to

$$\beta_{kt} = \frac{w_k f(\mathbf{y}_t|\mathbf{\Psi}_k)}{\sum_{k'} w_{k'} f(\mathbf{y}_t|\mathbf{\Psi}_{k'})}.$$

These probabilities are used during the maximization (M-step) to guide analytical parameter updates. In contrast, the direct optimization procedure only evaluates the log-likelihood function and updates parameters directly using an optimizer and automatic differentiation. Although the E-step is not required for direct numerical optimization, it is still used to compute posterior probabilities based on the estimated parameters.

Implementation details

We implemented mixture models with analytical maximization in NumPy and numerical optimization in PyTorch, where the latter also included the possibility for computing hidden Markov models (Supplementary Information). Recent advances in fast numerical gradient optimization frameworks including automatic differentiation in PyTorch has enabled fast non-linear model inference while only requiring the specification of an objective function, in this case the log-likelihood function. An advantage of direct numerical optimization when compared to the more conventional EM inference procedure is that the CACG and Gaussian models with low-rank-plus-diagonal reparameterization all require iterative fixed-point estimation for analytical maximization in the M-step. Conversely, direct numerical optimization requires choosing a learning rate and learning algorithm; Here we used the ADAM optimizer [67] with $\mathrm{LR} = 0.1$ in correspondence with adjacent studies [34,68].

Parameter constraints for numerical optimization were implemented through reparametrization of unconstrained variables to comply with the imposed constraints: We used a softmax function for the mixture weights $w_k = \frac{e^{\tilde{w}_k}}{\sum_{k'} e^{\tilde{w}_{k'}}}$ and a softplus function for the γ-parameter for Gaussian distributions to ensure non-negativity.

Mixture models, both using analytical and numerical optimization, were evaluated until a relative difference between the best and second-best iteration (divided by the best) in the latter 50 epochs fell below a tolerance level of $\mathrm{tol} = 10^{-10}$. While the EM algorithm typically guarantees monotone convergence, this is not the case for the low-rank-plus-diagonal reparameterizations, since these rely on iterative optimization procedures in the M-step. Analytical (EM) fixed-point updates for CACG and Gaussian distributions were evaluated until an absolute tolerance of 10^{-10}.

K-means equivalents

To provide a comparison to the probabilistic mixture models and to function as initialization to these, we utilized existing hard assignment K-means models for the projective hyperplane and introduced a complex-valued version of diametrical clustering [69] for the complex projective hyperplane.

Initialization schemes

Given the vast array of available models, gradually increasing model complexity may aid in ensuring model convergence to a suitable optimum. The K-means algorithms were initialized using their ++-versions, where an initial centroid is picked randomly and the remaining $K-1$ centroids iteratively picked based on a probability distribution proportional to the distance to existing centroids, ensuring that the starting points were spread far apart [70]. This scheme provided slightly better results than initializing centroids via the random uniform distribution (Supplementary Fig. S10). Probabilistic mixture models were initialized using centroids derived from K-means. The low-rank matrix $\mathbf{M}$ was initialized, given a partition of the data set $\mathbf{y}_t \in \mathcal{C}_k$, e.g., from K-means, as

$$\mathbf{M} = \mathbf{U}\left[\frac{1}{\epsilon}\left(\mathbf{\Sigma}_r - \epsilon\mathbf{I}\right)\right]^{1/2},$$

where $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\mathrm{H}} = \mathbf{Y}$ is the singular value decomposition and $\mathbf{Y} = [\mathbf{y}_t]_{\mathbf{y}\in\mathcal{C}_k}$ are the samples assigned to cluster k stacked horizontally. $\mathbf{\Sigma}_r$ is the diagonal matrix of singular values $[\sigma_{ll}]_{l=1}^{r}$, where all other values than the leading r singular values were set to zero, and $\epsilon = \frac{\sum_{l=r+1}^{p}\sigma_l}{p-r}$, assuming the number of assigned samples to cluster k is $n_k > p$. Mixture component weights $\mathbf{w}$ were initialized based on the initialization partition. For the Gaussian distributions, γ was initialized to the same value as ϵ.

We also investigated a scenario where high-rank mixtures were initialized by their lower-rank counterparts. For instance, given the parameter $\mathbf{M}_h$ of a rank $h < r$ model, the first h columns of $\mathbf{M}$ could be initialized as $\mathbf{M}_h$ and the remaining $r-h$ columns as the singular value decomposition on the residual matrix after projecting out the existing columns, i.e., $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\top} = \mathbf{Y} - \mathbf{M}_h\left(\mathbf{M}_h^{\mathrm{H}}\mathbf{M}_h\right)^{-1}\mathbf{M}_h^{\top}\mathbf{Y}$. Gradual initialization provided improved results (Supplementary Fig. S3).

Model evaluation

We evaluated our methods on both synthetic, phase-manipulated, and real task and resting-state fMRI data. To compare posterior probabilities to a known ground truth, we used normalized mutual information (NMI), which is invariant to permutation in clustering labels. Given two binary or weighted responsibility matrices $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ of size $K_1 \times N$ and $K_2 \times N$, respectively, which sum to one for each column, NMI is defined as [71]

$$\mathrm{NMI}(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2) = \frac{2\mathrm{MI}(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2)}{\mathrm{MI}(\boldsymbol{\beta}_1, \boldsymbol{\beta}_1) + \mathrm{MI}(\boldsymbol{\beta}_2, \boldsymbol{\beta}_2)},$$

$$\mathrm{MI}(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2) = \sum_{k_1,k_2} p(k_1,k_2) \log \frac{p(k_1,k_2)}{p(k_1)p(k_2)},$$

$$p(k_1,k_2) = \frac{1}{N} \sum_n \beta_{k_1,n}\beta_{k_2,n}.$$

Synthetic data was generated from two true three-dimensional components with consistent phase shifts and an added noise level $\sigma = \pi/16, \ldots, \pi/2$, $n = 1000$ samples for each component (Supporting Fig. S1). Models were fitted for each noise level separately and NMI evaluated against the true component labels.

To evaluate performance upon increasing model complexity (Supporting Fig S2), we used phase-controlled 116-dimensional resting-state data from one subject and one scan, i.e., $N = 1200$ data points corresponding to 15 minutes of data acquisition. First we divided both N and p into $K = 5$ groups, such that cluster 1 corresponded to the first 240 time points and the first 23 regions, cluster 2 corresponded to time points 241 to 480 and regions 24-46 and so on. Then we computed, for all regions and temporal windows, the Fourier spectrum and if the region and window corresponded to one of the clusters, the phase information for the relevant frequencies was set to a pre-specified constant phase $\alpha_j \in [0; 2\pi]$. If not, the phase information within the frequency band was randomized. We subsequently evaluated NMI between the posterior probability and one-hot encoded true labels.

Task fMRI scans were concatenated for train and test subjects into a full train and test set. Seven-component models were estimated on the full train set without access to labels, and subsequently the posterior probabilities were computed on the test set. Separate mixture models were estimated with component ranks $r \in [1, 5, 10, 25, 50, 100]$. Only the best of 10 estimates, evaluated using train set likelihood, were analyzed. We estimated NMI between the test set posterior probability and the one-hot-encoded true task labels. To compute a classification accuracy, we binarized the posterior at each sample to obtain an estimated "class" prediction, which we compared with true task labels, i.e., no additional classifier was trained. We first aligned each component with the task label that it corresponded to the most in the train set (average across all subjects) by solving a linear sum assignment problem and then evaluated classification accuracy on the test set. We also aggregated the posterior within each task scan to obtain an aggregated class prediction and computed a corresponding ensemble seven-class classification accuracy. Classification accuracies were tested between models using a Wilcoxon signed rank test Bonferroni-corrected for three comparisons.

Component covariance matrices were reduced from 116 parcels to 8 cortical and subcortical networks (Figs. 4F, S5-6). To assess which edges were specific to individual components, we first computed weighted complex-valued covariance matrices for each test subject and component using the test set posterior probabilities as weights. For each component, complex-valued covariance elements were compared to the corresponding elements in all other components using pairwise paired Hotelling's T2-test [38] on the test subjects (N=99 for task, N=100 for resting-state). For diagonal elements (self-connections), and the non-complex Gaussian model, we instead applied paired t-tests. This resulted in $p(p-1)/2 \cdot K(K-1)/2$ ($p = 8$ canonical networks) statistical tests per model for each of the three models considered (CACG, Gaussian, complex Gaussian). To summarize component specificity for a given network edge, we retained the maximum p-value across all $K-1$ component pairs involving that edge. The combined $K \cdot p(p-1)/2$ p-values were Bonferroni corrected for multiple comparisons.

On resting-state fMRI data we fitted models with varying model order K and component rank r. Each model was fitted 10 times and computational reproducibility evaluated using NMI between the posterior of all pairs of estimates. On the best estimate (evaluated using train set likelihood), we computed, for each test subject, the average sample log-likelihood and posterior entropy, computed as

$$H = -\sum_{i=1}^{N}\sum_{k=1}^{K} \beta_{ki} \log \beta_{ki},$$

where $N = 2400$ is the number of volumes per test subject.

Code and data availability

All processing and analysis scripts used in the present study, implemented primarily in Python, are made available at github.com/anders-s-olsen/PCMM. Human connectome project data are readily available online.

AUTHOR CONTRIBUTIONS

ASO and MM conceived the study and developed the theoretical framework. ASO performed the experiments, wrote the toolbox, and drafted the manuscript. AB aided in model development and PMF aided in analysis of results. MM and PMF supervised the project. All authors contributed to interpretation of results and reviewed and approved the final version of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

REFERENCES

1. Coombes, S. & Bressloff, P. C. *Bursting: The Genesis of Rhythm in the Nervous System*. (World Scientific, 2005).
2. Maeda, E., Robinson, H. & Kawana, A. The mechanisms of generation and propagation of synchronized bursting in developing networks of cortical neurons. *J. Neurosci.* **15**, 6834–6845 (1995).
3. Kirschstein, T. & Köhling, R. What is the Source of the EEG? *Clin. EEG Neurosci.* **40**, 146–149 (2009).
4. Bola, M. & Sabel, B. A. Dynamic reorganization of brain functional networks during cognition. *NeuroImage* **114**, 398–413 (2015).
5. Uhlhaas, P. *et al.* Neural synchrony in cortical networks: history, concept and current status. *Front. Integr. Neurosci.* **3**, (2009).
6. Womelsdorf, T. *et al.* Modulation of Neuronal Interactions Through Neuronal Synchronization. *Science* **316**, 1609–1612 (2007).
7. Biswal, B., Zerrin Yetkin, F., Haughton, V. M. & Hyde, J. S. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magn. Reson. Med.* **34**, 537–541 (1995).
8. Bastos, A. M. & Schoffelen, J. M. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. *Front. Syst. Neurosci.* **9**, (2016).

9. Wang, H. E. *et al.* A systematic framework for functional connectivity measures. *Front. Neurosci.* **8**, (2014).
10. Allen, E. A. *et al.* Tracking whole-brain connectivity dynamics in the resting state. *Cereb. Cortex* **24**, 663–676 (2014).
11. Leonardi, N. & Van De Ville, D. On spurious and real fluctuations of dynamic functional connectivity during rest. *NeuroImage* **104**, 430–436 (2015).
12. Nielsen, S. F. V., Madsen, K. H., Schmidt, M. N. & Morup, M. Modeling dynamic functional connectivity using a wishart mixture model. *2017 Int. Workshop Pattern Recognit. Neuroimaging PRNI 2017* https://doi.org/10.1109/PRNI.2017.7981505 (2017) doi:10.1109/PRNI.2017.7981505.
13. Preti, M. G., Bolton, T. A. & Van De Ville, D. The dynamic functional connectome: State-of-the-art and perspectives. *NeuroImage* **160**, 41–54 (2017).
14. Choe, A. S. *et al.* Comparing test-retest reliability of dynamic functional connectivity methods. *NeuroImage* **158**, 155–175 (2017).
15. Hindriks, R. *et al.* Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? *NeuroImage* **127**, 242–256 (2016).
16. Laumann, T. O. *et al.* On the Stability of BOLD fMRI Correlations. *Cereb. Cortex* **27**, 4719–4732 (2017).
17. Liégeois, R., Laumann, T. O., Snyder, A. Z., Zhou, J. & Yeo, B. T. T. Interpreting temporal fluctuations in resting-state functional connectivity MRI. *NeuroImage* **163**, 437–455 (2017).
18. Laumann, T. O., Snyder, A. Z. & Gratton, C. Challenges in the measurement and interpretation of dynamic functional connectivity. *Imaging Neurosci.* https://doi.org/10.1162/imag_a_00366 (2024) doi:10.1162/imag_a_00366.
19. Vidaurre, D., Smith, S. M. & Woolrich, M. W. Brain network dynamics are hierarchically organized in time. *Proc. Natl. Acad. Sci. U. S. A.* **114**, 12827–12832 (2017).
20. Chen, J. E., Rubinov, M. & Chang, C. Methods and Considerations for Dynamic Analysis of

Functional MR Imaging Data. *Neuroimaging Clin. N. Am.* **27**, 547–560 (2017).

21. Glerean, E., Salmi, J., Lahnakoski, J. M., Jääskeläinen, I. P. & Sams, M. Functional Magnetic Resonance Imaging Phase Synchronization as a Measure of Dynamic Functional Connectivity. *Brain Connect.* **2**, 91–101 (2012).
22. Cabral, J. *et al.* Cognitive performance in healthy older adults relates to spontaneous switching between states of functional connectivity during rest. *Sci. Rep.* **7**, 1–13 (2017).
23. Honari, H., Choe, A. S. & Lindquist, M. A. Evaluating phase synchronization methods in fMRI: A comparison study and new approaches. *NeuroImage* **228**, 117704 (2021).
24. Kringelbach, M. L. *et al.* Dynamic coupling of whole-brain neuronal and neurotransmitter systems. *Proc. Natl. Acad. Sci.* **117**, 9566–9576 (2020).
25. Zarghami, T. S., Hossein-Zadeh, G.-A. & Bahrami, F. Deep Temporal Organization of fMRI Phase Synchrony Modes Promotes Large-Scale Disconnection in Schizophrenia. *Front. Neurosci.* **14**, (2020).
26. Nolte, G. *et al.* Identifying true brain interaction from EEG data using the imaginary part of coherency. *Clin. Neurophysiol.* **115**, 2292–2307 (2004).
27. Kent, J. T. Data analysis for shapes and images. *J. Stat. Plan. Inference* **57**, 181–193 (1997).
28. Tyler, D. E. Statistical Analysis for the Angular Central Gaussian Distribution on the Sphere. *Biometrika* **74**, 579 (1987).
29. Shine, J. M. *et al.* The Low-Dimensional Neural Architecture of Cognitive Complexity Is Related to Activity in Medial Thalamic Nuclei. *Neuron* **104**, 849-855.e3 (2019).
30. Rué-Queralt, J. *et al.* Decoding brain states on the intrinsic manifold of human brain dynamics across wakefulness and sleep. *Commun. Biol.* **4**, 854 (2021).
31. Van Essen, D. C. *et al.* The WU-Minn Human Connectome Project: an overview. *NeuroImage* **80**, 62–79 (2013).
32. Schaefer, A. *et al.* Local-Global Parcellation of the Human Cerebral Cortex from Intrinsic

Functional Connectivity MRI. *Cereb. Cortex* **28**, 3095–3114 (2018).

33. Tian, Y., Margulies, D. S., Breakspear, M. & Zalesky, A. Topographic organization of the human subcortex unveiled with functional connectivity gradients. *Nat. Neurosci.* **23**, 1421–1432 (2020).
34. Olsen, A. S., Ortvald, E., Madsen, K. H., Schmidt, M. N. & Mørup, M. Angular Central Gaussian and Watson Mixture Models for Assessing Dynamic Functional Brain Connectivity During a Motor Task. in *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)* 1–5 (2023). doi:10.1109/ICASSPW59220.2023.10193021.
35. Zhang, Y., Tetrel, L., Thirion, B. & Bellec, P. Functional annotation of human cognitive states using deep graph convolution. *NeuroImage* **231**, 117847 (2021).
36. Rastegarnia, S., St-Laurent, M., DuPre, E., Pinsard, B. & Bellec, P. Brain decoding of the Human Connectome Project tasks in a dense individual fMRI dataset. *NeuroImage* **283**, 120395 (2023).
37. Thomas Yeo, B. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J. Neurophysiol.* **106**, 1125–1165 (2011).
38. Hotelling, H. The Generalization of Student's Ratio. *Ann. Math. Stat.* **2**, 360–378 (1931).
39. Hannum, A., Lopez, M. A., Blanco, S. A. & Betzel, R. F. High-accuracy machine learning techniques for functional connectome fingerprinting and cognitive state decoding. *Hum. Brain Mapp.* **44**, 5294–5308 (2023).
40. Gupta, S., Lim, M. & Rajapakse, J. C. Decoding task specific and task general functional architectures of the brain. *Hum. Brain Mapp.* **43**, 2801–2816 (2022).
41. Wang, X. *et al.* Decoding and mapping task states of the human brain via deep learning. *Hum. Brain Mapp.* **41**, 1505–1519 (2019).
42. Olsen, A. S. *et al.* Psilocybin modulation of time-varying functional connectivity is associated with plasma psilocin and subjective effects. *NeuroImage* **264**, 119716 (2022).

43. Castro, P. *et al.* Dynamical structure-function correlations provide robust and generalizable signatures of consciousness in humans. *Commun. Biol.* **7**, 1–12 (2024).
44. de Alteriis, G. *et al.* EiDA: A lossless approach for dynamic functional connectivity; application to fMRI data of a model of ageing. *Imaging Neurosci.* **2**, 1–22 (2024).
45. Demertzi, A. *et al.* Human consciousness is supported by dynamic complex patterns of brain signal coordination. *Sci. Adv.* **5**, eaat7603 (2019).
46. Lord, L. D. *et al.* Dynamical exploration of the repertoire of brain networks at rest is modulated by psilocybin. *NeuroImage* **199**, 127–142 (2019).
47. Olsen, A. S. *et al.* Psilocybin modulation of time-varying functional connectivity is associated with plasma psilocin and subjective effects. *NeuroImage* 119716 (2022) doi:10.1016/J.NEUROIMAGE.2022.119716.
48. Vohryzek, J., Deco, G., Cessac, B., Kringelbach, M. L. & Cabral, J. Ghost Attractors in Spontaneous Brain Activity: Recurrent Excursions Into Functionally-Relevant BOLD Phase-Locking States. *Front. Syst. Neurosci.* **14**, 20 (2020).
49. Pedersen, M., Omidvarnia, A., Zalesky, A. & Jackson, G. D. On the relationship between instantaneous phase synchrony and correlation-based sliding windows for time-resolved fMRI connectivity analysis. *NeuroImage* **181**, 85–94 (2018).
50. Pikovsky, A., Rosenblum, M. & Kurths, J. *Synchronization: A Universal Concept in Nonlinear Sciences*. (Cambridge University Press, Cambridge, 2001). doi:10.1017/CBO9780511755743.
51. Honari, H. & Lindquist, M. A. Mode decomposition-based time-varying phase synchronization for fMRI. *NeuroImage* **261**, 119519 (2022).
52. Ahrends, C. *et al.* Data and model considerations for estimating time-varying functional connectivity in fMRI. *NeuroImage* **252**, 119026 (2022).
53. Liégeois, R. *et al.* Resting brain dynamics at different timescales capture distinct aspects of human behavior. *Nat. Commun.* **10**, 2317 (2019).

54. Stevner, A. B. A. *et al.* Discovery of key whole-brain transitions and dynamics during human wakefulness and non-REM sleep. *Nat. Commun.* **10**, 1035 (2019).
55. Vidaurre, D. *et al.* Spectrally resolved fast transient brain states in electrophysiological data. *NeuroImage* **126**, 81–95 (2016).
56. Huang, N. E. *et al.* The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. Lond. Ser. Math. Phys. Eng. Sci.* **454**, 903–995 (1998).
57. Elam, J. S. *et al.* The Human Connectome Project: A retrospective. *NeuroImage* **244**, 118543 (2021).
58. Barch, D. M. *et al.* Function in the Human Connectome: Task-fMRI and Individual Differences in Behavior. *NeuroImage* **80**, 169 (2013).
59. Smith, S. M. *et al.* Resting-state fMRI in the Human Connectome Project. *NeuroImage* **80**, 144–168 (2013).
60. Glasser, M. F. *et al.* The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* **80**, 105–124 (2013).
61. Glasser, M. F. *et al.* Using temporal ICA to selectively remove global noise while preserving global signal in functional MRI data. *NeuroImage* **181**, 692–717 (2018).
62. Feldman, M. Hilbert transform in vibration analysis. *Mech. Syst. Signal Process.* **25**, 735–802 (2011).
63. Sun, J. & Small, M. Unified framework for detecting phase synchronization in coupled time series. *Phys. Rev. E* **80**, 046219 (2009).
64. Boashash, B. Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals. *Proc. IEEE* **80**, 520–538 (1992).
65. Figueroa, C. A. *et al.* Altered ability to access a clinically relevant control network in patients remitted from major depressive disorder. *Hum. Brain Mapp.* **40**, 2771–2786 (2019).
66. Mardia, K. V. & Jupp, P. E. *Directional Statistics*. *Directional Statistics* (John Wiley and Sons

Ltd., 1999). doi:10.1002/9780470316979.

67. Kingma, D. P. & Ba, J. L. Adam: A Method for Stochastic Optimization. *3rd Int. Conf. Learn. Represent. ICLR 2015 - Conf. Track Proc.* https://doi.org/10.48550/arxiv.1412.6980 (2014) doi:10.48550/arxiv.1412.6980.
68. Olsen, A. S., Nielsen, J. D. & Mørup, M. Coupled Generator Decomposition for Fusion of Electro- and Magnetoencephalography Data. in *2024 32nd European Signal Processing Conference (EUSIPCO)* 1357–1361 (2024).
69. Dhillon, I. S., Marcotte, E. M. & Roshan, U. Diametrical clustering for identifying anti-correlated gene clusters. *Bioinformatics* **19**, 1612–1619 (2003).
70. Arthur, D. & Vassilvitskii, S. K-means++: The advantages of careful seeding. in *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms* vols 07-09-Janu 1027–1035 (2007).
71. Hinrich, J. L. *et al.* Archetypal Analysis for Modeling Multisubject fMRI Data. *IEEE J. Sel. Top. Signal Process.* **10**, 1160–1171 (2016).